arXiv:1605.07879v1 [math.GR] 25 May 2016

CUBULATING MAPPING TORI OF POLYNOMIAL GROWTH FREE GROUP AUTOMORPHISMS

MARK F. HAGEN AND DANIEL T. WISE

ABSTRACT. Let $\Phi : F \to F$ be a polynomially-growing automorphism of a finite-rank free group F. Then $G = F \rtimes_\Phi \mathbb{Z}$ acts freely on a CAT(0) cube complex.

1. INTRODUCTION

The goal of this paper is to prove:

Theorem A. *Let F be a finite-rank free group and let $\Phi : F \to F$ be a polynomially-growing automorphism. Then $G = F \rtimes_\Phi \mathbb{Z}$ acts freely on a CAT(0) cube complex.*

We emphasize that the action of G on the cube complex above necessarily has exotic properties: it is not in general metrically proper, and the cube complex is not locally finite and usually contains infinite cubes, as we discuss in more detail below.

Theorem A applies to natural examples coming from a diverse and well-studied class of groups, mapping tori of free-group automorphisms. There has been a wave of recent interest in this topic, whose broad appeal is illustrated by the many recent results; see e.g. [1, 2, 3, 10, 11, 15, 16, 17, 18, 19, 20, 30, 21, 31, 32, 33, 34, 35, 37].

Relative train track maps provide a fundamental tool for studying free-by-cyclic groups (see [9, 7, 8]) and classifies them according to the dynamics of the monodromy. After replacing Φ with a power (thus replacing G with a finite-index subgroup), Φ is represented by a so-called *improved relative train track map* in the sense of [7]. In this paper, we focus on groups of the form $G = F \rtimes_\Phi \langle t \rangle$ where Φ is *polynomially growing* in the sense that there exists a polynomial h so that $|\Phi^n(f)| \leq h(n)|f|$ for each $f \in F$. We call such a G a *polynomial free-by-$\mathbb{Z}$ group*. The nature of the improved relative train track in the polynomial case shows that G splits as an iterated HNN extension with cyclic edge groups at each stage. Specifically, $G = G_n$, where $G_{i+1} = \langle G_i, e_i \mid e_i t e_i^{-1} = p_i t \rangle$ where $p_i \in G_i$. From the viewpoint of the relative train track map, this sequence would terminate with $G_0 = \langle t \rangle$. However, in our text we will find that it suffices, and is useful, to consider the case where G_0 is the centralizer of t.

This sequence of splittings has been used to prove various results illuminating the geometry of polynomial free-by-$\mathbb{Z}$ groups. For example, in [35], Macura studied the divergence function of these groups, a helpful quasi-isometry invariant, showing that the divergence of G is $\sim x^{r+1}$, where r is

Date: May 26, 2016.

2010 *Mathematics Subject Classification.* Primary: 20F65; Secondary: 57M20.

Key words and phrases. free-by-cyclic group, CAT(0) cube complex, relative train track map, polynomial-growth automorphism.

Hagen was supported by NSF Grant Number 1045119 and by the EPSRC grant of Henry Wilton.
Wise was supported by NSERC, the IHP, and a Lady Davis Fellowship at the Technion.

the order of the (polynomial) growth of Φ. In fact, it is implicit in [35] that G is *thick of order* r in the sense of [5].

Several other theorems on the geometry of free-by-$\mathbb{Z}$ groups illustrate the presence or absence of various nonpositive curvature properties. A free-by-$\mathbb{Z}$ group is hyperbolic when it contains no $\mathbb{Z}^2$ subgroup, i.e. when F has no Φ–periodic nontrivial conjugacy class [6, 12]. However, in the polynomial case, G is never hyperbolic (unless $G \cong \mathbb{Z}$ which happens in the degenerate case where F is trivial). In fact, polynomial free-by-$\mathbb{Z}$ groups need not be CAT(0), or even admit a proper semisimple action on a CAT(0) space, as shown by the following example of Gersten [22]:

$$\langle a, b, c, t \mid a^t = a, b^t = ba, c^t = ca^2 \rangle$$

Nonetheless, polynomial free-by-$\mathbb{Z}$ groups do enjoy properties reminiscent of nonpositive curvature: Macura showed that they satisfy a quadratic isoperimetric inequality [34], and this was generalized by Bridson-Groves to arbitrary free-by-$\mathbb{Z}$ groups [11]. Both results use improved relative train track maps.

The story of free-by-$\mathbb{Z}$ groups acting on CAT(0) cube complexes is arduous. If $G = F \rtimes_\Phi \mathbb{Z}$ is hyperbolic then G acts freely and cocompactly on a CAT(0) cube complex [24, 25]. When F has rank 2, the automorphism is represented by a surface group automorphism, and virtual cocompactness holds by [40] and [23] since the mapping torus either only has hyperbolic pieces or is a graph manifold. Moreover, Button-Kropholler [14] have gone further and shown directly that G is the fundamental group of a compact nonpositively curved 2-dimensional cube complex. However, in view of Haglund's semisimplicity theorem [26], Gersten's polynomial free-by-$\mathbb{Z}$ group cannot act freely on a proper CAT(0) cube complex.

In addition to being free-by-$\mathbb{Z}$, Gersten's group is *tubular*, i.e. it splits as a graph of groups with $\mathbb{Z}^2$ vertex groups and cyclic edge groups. There are many tubular groups that act freely on CAT(0) cube complexes, even though they cannot act metrically properly. A necessary and sufficient condition for acting freely was provided in [38], and in particular, Gersten's group satisfies this condition. Recently, Button characterized the free-by-$\mathbb{Z}$ groups that are tubular, and applied the condition from [38] to see that every tubular free-by-$\mathbb{Z}$ group acts freely on a CAT(0) cube complex [13]. Since tubular groups are thick of order ≤ 1 relative to the $\mathbb{Z}^2$ vertex groups, results of [35, 4] show that if $G \cong F \rtimes_\Phi \mathbb{Z}$ is tubular, then Φ is linearly-growing.

Theorem A is motivated by many of the above results. Most directly, it generalizes the tubular/linear-growth case to arbitrary polynomial free-by-$\mathbb{Z}$ groups and also provides a non-hyperbolic counterpart to the cubulation of hyperbolic free-by-$\mathbb{Z}$ groups. It also broadens the agenda of [38]: instead of studying a single splitting of a group as a graph of groups with cyclic edge groups, we are instead studying a sequence of iterated cyclic HNN extensions, where the edge groups at each stage can involve stable letters from previous stages. Thus, although the initial framework of the proof of Theorem A is superficially similar to that used for tubular groups, our situation is significantly more complicated.

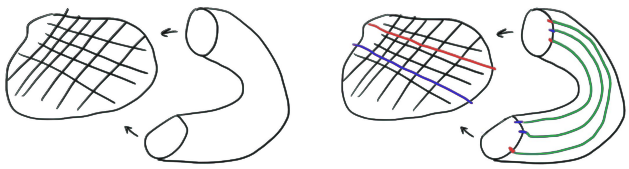

FIGURE 1

As shown at left in Figure 1, the main object in this paper is a space D consisting of a nonpositively curved cube complex with a cylinder attached along its boundary circles, which we may assume are local geodesics. The fundamental group of D is G_n, and the cube complex has fundamental group G_{n-1}. Attaching the cylinder performs the HNN extension. If the left attaching circle and the right attaching circle have the same number of intersection points with hyperplanes, then as shown in Figure 1 at right, we could join these points in pairs, by arcs in the cylinder, to create immersed walls. The result would yield a wallspace on $\widetilde{D}$, and G_n would act freely on the dual cube complex. However, in our situation, as D is initially presented to us, the attaching circles are rarely the same length. Most of the proof is a juggling act to remedy this.

The results of this paper contribute to the problem of cubulating a cyclic HNN extension of a cubulated group. When D is word-hyperbolic, this problem has a positive solution [29]. However, this is not the case in general. For instance, the Baumslag-Solitar group $BS(1,2) \cong \langle e, a \mid eae^{-1} = a^2 \rangle$ cannot act freely on a CAT(0) cube complex [26], and even assuming that there are no bad Baumslag-Solitar subgroups, there is an even an example of an HNN extension of a nonpositively curved square complex with a special double cover which cannot act freely on a CAT(0) cube complex [29, Example 8.7].

1.1. **Summary of the proof.** Let $G = F \rtimes_\Phi \langle t \rangle$. We first reduce to the case where Φ is represented by an improved relative train track map. As discussed above, this enables us to decompose G as an iterated HNN extension $G = G_n \cong \langle G_{n-1}, e_n \mid e_n t e_n^{-1} = p_n t \rangle$, where G_{n-1} is a polynomial free-by-cyclic group and $p_n, t \in G_{n-1}$. The stable letter e_n comes from the top stratum of the relative train track representative. With some more care, we can also arrange for G_0 to be the centralizer of t in G, which is crucial throughout the proof. We then induct on n.

1.1.1. *Cubulating $F \times \mathbb{Z}$, and finite-dimensionality.* In the base case, $n = 0$, we must produce a free action of $F \times \langle t \rangle$ on a CAT(0) cube complex, but an additional consideration means that we cannot use the obvious one: for each of the "future edge groups" $\langle p_i t \rangle$ that are already visible at this stage (i.e. for which p_i centralizes t), we must ensure that the translation length of $p_i t$ on our cube complex is the same as that of t. Indeed, we will later attach a cylinder to our cube complex, whose boundary circles correspond to $\langle t \rangle$ and $\langle p_i t \rangle$. In order to pair the hyperplane points on one of these circles with those on the other, we need these translation lengths to coincide.

Combining Theorem 4.1 and Remark 4.4 yields the following, which we feel to be of independent interest:

Theorem B. *Let $\mathcal{P}$ be a finite set of nontrivial elements of the finitely-generated free group F. Then there exists a metrically proper action of $F \times \langle t \rangle$ on a CAT(0) cube complex C so that $\|pt\| = \|t\|$ for all $p \in \mathcal{P}$.*

An artifact of our construction is that C is always infinite-dimensional. In some cases, this may not be necessary (see below).

Theorem B is proved in two stages: first, we construct a (non-free) action of F on a CAT(0) cube complex $C_{\copyright}$ so that the nontrivial 0–cube stabilizers are exactly the conjugates of the maximal cyclic subgroups containing the various $p \in \mathcal{P}$. Next, we construct a (non-free) action of $F \times \langle t \rangle$ on a CAT(0) cube complex $C_\bullet$ on which $t, pt, p \in \mathcal{P}$ all have the same positive translation length. The resulting action of $F \times \langle t \rangle$ on $C_{\copyright} \times C_\bullet$ has the desired properties.

The cube complex we eventually construct in the proof of Theorem A is always infinite-dimensional because the cube complex produced by Theorem B is, but this appears to be an artifact of the construction of $C_{\copyright}$; see Problem 1 below.

1.1.2. *The inductive step.* The inductive part of the proof is carried out in Theorem 5.2. Recall that we have a sequence $G_0 = Z_t(G)$, $G_{i+1} = \langle G_i, e_{i+1} \mid e_{i+1} t e_{i+1}^{-1} = p_i t \rangle$, $0 \leq i \leq n-1$, with $G = G_n$. In the inductive step, we assume that G_i acts freely on a CAT(0) cube complex C_0, and for each element $p \in G_i$ which is conjugate in G to some element of $\{p_0, \ldots, p_{n-1}\}$, we assume that the translation length of pt on C_0 is the same as that of t. (In the statement of Theorem 5.2, there are also various other hypotheses. These reflect both the fact that we work in the category of cyclic HNN extensions with specified properties that hold for free-by-$\mathbb{Z}$ groups and also some extra technical challenges that are taken care of inductively.)

For each $q \in G_{i+1} - G_i$ which is conjugate to some element of $\{p_0, \ldots, p_{n-1}\}$, we need to ensure that the translation length of qt on the cube complex we eventually construct coincides with that of t (and with each of the pt above).

We work in the space D, shown in Figure 1, obtained from $G_i \backslash C_0$ by attaching a cylinder $B \times [-1, 1]$ along circles $B \times \{\pm 1\}$ corresponding to $t, p_i t$. Note that $\pi_1 D \cong G_{i+1}$. The first step is to construct immersed walls in D, each of which is a finite graph of spaces whose vertex spaces are hyperplanes in $G_i \backslash C_0$ and whose edge spaces are arcs in $B \times [-1, 1]$ joining hyperplane points of $B \times \{-1\}$ to hyperplane points of $B \times \{1\}$. We do this in such a way that for all $pt, qt, p't, q't$ as above, if some lift $\widetilde{W}$ of some immersed wall $W \to D$ cuts the axis of $a \in \{pt, qt, p't, q't\}$ and that of $b \in \{pt, qt, p't, q't\}$, then the translates of $\widetilde{W}$ by high powers of b are disjoint. Although we think that this choice of immersed walls can be made constructively, the combinatorics are tricky and instead we use a probabilistic argument, in Section 5.2, to show that choosing the immersed walls $W \to D$ randomly yields a collection with the desired property. We also include the edge-spaces as walls. The above cutting property persists when we pass to the action of G_i on the dual cube complex $\widetilde{X}$.

The next step is to regard $\widetilde{X}$ as a tree of spaces, with underlying tree corresponding to the original HNN extension. Each pt is elliptic and each qt hyperbolic on the tree. To each qt we can associate a convex subspace $\widetilde{Y} \subset \widetilde{X}$ which is the preimage of the qt–axis in $\widetilde{X}$ under the projection to the underlying tree. The inclusion $\widetilde{Y} \to \widetilde{X}$ descends to a local isometry $\widetilde{Y} \to G_{i+1} \backslash \widetilde{X}$, and we can regard G_{i+1} as the fundamental group of the space obtained from $G_{i+1} \backslash \widetilde{X}$ by coning off the various $\widetilde{Y}$. This gives G_{i+1} a cubical presentation which satisfies a simplified form of the $B(6)$ condition from [40], once we endow each $\widetilde{Y}$ with an alternative wallspace structure. The cubical small-cancellation technology from [40] then produces an action of G_{i+1} on a CAT(0) cube complex where the walls that formerly cut both qt and $q't$ have been "turned away" from $q't$ and now cut qt but not $q't$ (because of the particular choice of wallspace structure on $\widetilde{Y}$). Another application of this method allows us to arrange for various walls to cut the various hyperbolic qt while avoiding the elliptic pt. Once we are in this situation, the proof can be completed by equivariantly subdividing the cube complex, so each hyperplane is replaced with some number of parallel copies, equalizing the translation lengths of the pt, qt and t.

1.2. **Problems.** We conclude with problems suggested by the proof of Theorem A.

1.2.1. *Metrically proper actions.* As we are dealing with CAT(0) cube complexes that may not be locally finite or finite-dimensional, we must distinguish between free actions and metrically proper actions. In view of Theorem A, it is reasonable to ask:

Problem 1. Which polynomial free-by-cyclic groups act metrically properly on CAT(0) cube complexes, and, slightly more strongly, which such groups act freely on proper CAT(0) cube complexes? For which polynomially-growing Φ can one produce a free action of G on a finite-dimensional

CAT(0) cube complex? Additionally, we wonder if there is a characterization of the polynomially-growing Φ for which $F \rtimes_\Phi \mathbb{Z}$ (virtually) acts freely and cocompactly on a CAT(0) cube complex.

For tubular groups, a criterion for finite-dimensional cubulation was given by Woodhouse [41], which is part of an ongoing project showing that tubular groups acting freely on locally finite CAT(0) cube complexes are virtually special. We wonder about the implications between the following statements for a polynomial free-by-$\mathbb{Z}$ group G:

- G acts metrically properly on a CAT(0) cube complex;
- G acts freely on a locally finite CAT(0) cube complex;
- G acts freely on a uniformly locally finite CAT(0) cube complex;
- G is virtually special.

The class of tubular free-by-$\mathbb{Z}$ groups shows that among G with Φ linearly-growing one can already find examples where no metrically proper action exists. Accordingly, an answer to Problem 1 must involve more subtle properties of Φ than the growth. We suspect it has to do with the nature of the splitting of G as an iterated cyclic HNN extension $G_n = \langle G_{n-1}, e_n \mid e_n t e_n^{-1} = p_n t \rangle$. Intuitively, in the situation where the rate of polynomial growth is n, i.e. each $p_i \in G_{i-1}$ involves the stable letter in the splitting of G_{i-1}, we are in a "hyperbolic-like" situation where it is possible that there is an easier cubulation, possibly giving rise to a nicer cube complex and a metrically proper action.

The cubulations should be forced to be nasty if the edge-groups $\langle p_i t \rangle$ appear early in the hierarchy relative to i. Accordingly, consider the following situation: $G_0 \cong F_0 \times \langle t \rangle$ and $G_{i+1} = \langle G_t, e_{i+1} \mid e_{i+1} t e_{i+1}^{-1} = f_i t \rangle$ for $0 \leq i \leq n$, where each $f_i \in F_0$. Theorem B allows one to easily construct a free action on a CAT(0) cube complex in this case. For the purposes of obtaining a metrically proper action of $G = G_n$ on a cube complex, is this in some sense a worst-case scenario? For which choices of $\{f_i\}$ can one produce such an action?

1.2.2. *Arbitrary mapping tori.* Finally, the methodology used here is totally unrelated to the technique used to cubulate mapping tori of hyperbolic free group automorphisms in [25] [24]. However, we hope the methods will eventually be combined to resolve:

Problem 2. Show that every free-by-$\mathbb{Z}$ group acts freely on a CAT(0) cube complex.

We do not expect that a free action on a locally infinite CAT(0) cube complex will shed a large amount of light on the structure of the free-by-$\mathbb{Z}$ group, although there are some interesting consequences such as undistortedness of maximal finitely generated free abelian subgroups [42]. It is nonetheless intriguing that groups in this natural class admit such actions.

1.3. **Organization of the paper.** Section 2 contains background on relative train track maps, actions on CAT(0) cube complexes, cubical small-cancellation theory, and a useful lemma about actions on trees with $\mathbb{Z}$ edge-stabilizers. In Section 3 we prove Theorem A. This proof is inductive: the base case is Theorem 4.1, which is the subject of Section 4. The inductive step uses Theorem 5.2.

2. Background

2.1. **Relative train track maps for polynomial automorphisms.** Throughout, F is a finitely-generated free group and $\Phi \in \mathrm{Aut}(F)$.

Definition 2.1 (Polynomial growth)**.** The automorphism Φ is *polynomially growing* if for each $f \in F$, the word length $|\Phi^k(f)|$ grows polynomially as $k \to \infty$.

We set up our proof of Theorem A using a result of [7] saying that each element of $\mathrm{Aut}(F)$ has a nonzero power represented by an *improved relative train track map*, a homotopy equivalence of graphs with numerous useful properties. The full statement can be found in [7, Theorem 5.1.5], but we need just the following:

Proposition 2.2 (Relative train tracks for polynomially-growing automorphisms)**.** *Let F be a finitely-generated free group and let $\Phi : F \to F$ be a polynomially-growing automorphism. Then there exists $k > 0$, a connected graph Γ, and a homotopy equivalence $f : \Gamma \to \Gamma$, representing Φ^k, that sends vertices to vertices and has the following properties:*

1. *There is a sequence of f–invariant subgraphs $\Gamma^0 = \Gamma_0 \subset \Gamma_1 \ldots \subset \Gamma_r = \Gamma$ such that $\Gamma_i - \mathrm{Int}(\Gamma_{i-1})$ consists of a single edge E_i for $1 \leq i \leq r$;*
2. *$f(v) \in \Gamma^0$ for each $v \in \Gamma^0$;*
3. *$f(E_i) = E_i P_{i-1}$, where P_{i-1} is a closed combinatorial path in Γ_{i-1} whose basepoint is fixed by f. Thus $f(v) = v$ for all $v \in \Gamma^0$;*

We need to impose extra constraints on the filtration:

Proposition 2.3 (Cyclic hierarchy with Z_t at the bottom)**.** *Let F be a finitely-generated free group and let $\Phi \in \mathrm{Aut}(F)$ be polynomially growing with an IRTTM. Let $G = F \rtimes_\Phi \mathbb{Z}$, so*

$$G \;\cong\; \langle\, F, t \;\mid\; \{tft^{-1} = \Phi(f) : f \in F\} \,\rangle,$$

and let $Z_t \leq G$ denote the centralizer of t. Then G is virtually contained in a group $G' \cong F' \rtimes_{\Phi'} \langle t \rangle$ for some $\Phi' \in \mathrm{Aut}(F')$, such that there exists a sequence $Z_t = G'_0 \leq \cdots \leq G'_n = G'$ such that:

1. *$G'_i \cong F'_i \rtimes_{\Phi'} \langle t \rangle$, where $F'_i \leq F'$, for $0 \leq i \leq n$;*
2. *$G'_{i+1} = \langle G'_i, e_{i+1} \mid e_{i+1}^{-1} t e_{i+1} = p_{i+1} t \rangle$, where $p_{i+1} \in F'_i$ and $e_{i+1} \in F'_{i+1} - F'_i$, for $0 \leq i \leq n-1$.*

Moreover, Φ' has an improved relative train track representative $\hat{f} : \widehat{\Gamma} \to \widehat{\Gamma}$ corresponding to a filtration $\widehat{\Gamma}_0 \subset \ldots \subset \widehat{\Gamma}_n$ of $\widehat{\Gamma}$, where $\pi_1 \widehat{\Gamma}_0 \times \langle t \rangle \cong Z_t$ and $\mathrm{Cl}(\widehat{\Gamma}_i - \widehat{\Gamma}_{i-1})$ consists of a single edge E_i for $1 \leq i \leq n$.

Proof. Let $f : \Gamma \to \Gamma$ be the improved relative train track map from Proposition 2.2. Let $\overline{\Gamma}$ be formed from Γ by identifying all of the vertices, and let $\bar{f} : \overline{\Gamma} \to \Gamma$ be the corresponding quotient of f, which continues to satisfy the conclusion of Proposition 2.2. Let $\bar{t}$ be the edge of the mapping torus of $\bar{f}$ corresponding to the stable letter and joining the unique vertex $\bar{v}$ of $\overline{\Gamma}$ to itself. Consider the group $\overline{G} = \pi_1 \overline{\Gamma} \rtimes_{\overline{\Phi}} \langle \bar{t} \rangle$ and let $N_{\bar{t}}$ denote the normalizer of $\bar{t}$. Observe that $N_{\bar{t}}$ is separable, being the normalizer of an element of the residually finite group G (indeed, if $x\bar{t}^n x^{-1} t^{\pm n} \neq 1$, then for any finite quotient $\overline{G} \to Q$ in which $x\bar{t}^n x^{-1} t^{\pm n}$ survives, the image of x is not in the image of $N_{\bar{t}}$). Hence there is a finite characteristic cover $\widehat{\Gamma} \to \overline{\Gamma}$ so that $N_{\bar{t}}$ is represented by a subcomplex of the mapping torus of the lift $\hat{f} : \widehat{\Gamma} \to \widehat{\Gamma}$ of $\bar{f}$. By replacing $\hat{f}$ with a power, we can assume that $\hat{f}(v) = v$ for each $v \in \widehat{\Gamma}^0$. Moreover, $\hat{f}$ satisfies the conclusion of Proposition 2.2.

Form $\ddot{\Gamma}$ from $\widehat{\Gamma}$ by identifying all the vertices and consider $G' = \pi_1 \ddot{\Gamma} \rtimes_{\Phi'} \langle \ddot{t} \rangle$, where Φ' is represented by the quotient $\ddot{f} : \ddot{\Gamma} \to \ddot{\Gamma}$ of $\hat{f}$. Let Λ be the connected subgraph of $\ddot{\Gamma}$ consisting of all edges e such that $\ddot{f}(e) = e$. The map $\ddot{f}$ continues to satisfy the conclusion of Proposition 2.2, and we may assume that the edges of Λ constitute the lowest strata in the filtration of $\ddot{\Gamma}$ from Proposition 2.2. □

2.2. **Cubical translation length.** Let C be a CAT(0) cube complex. Then $g \in \mathrm{Aut}(C)$ is *(combinatorially) elliptic* if g fixes a 0–cube and *(combinatorially) hyperbolic* if there is a bi-infinite combinatorial geodesic $\widetilde{A}$ in C so that g stabilizes $\widetilde{A}$ in the following way: there exists $\tau \in \mathbb{Z}$ so

that, regarding $\widetilde{A}$ as an isometric embedding $\widetilde{A} : \mathbb{R} \to C^1$, we have $g\widetilde{A}(t) = \widetilde{A}(t + \tau)$ for all $t \in \mathbb{R}$. By replacing C by a *cubical subdivision* – i.e. by subdividing C so that the hyperplanes become subcomplexes and each hyperplane is replaced by 2 parallel hyperplanes – we can assume that any fixed group $G \leq \mathrm{Aut}(C)$ has the property that each $g \in G$ is either elliptic or hyperbolic [26]. Indeed, it is shown in [26] that the action of G on C has this property provided G acts *without inversions in hyperplanes* in the sense that whenever $g \in G$ stabilizes a hyperplane, it stabilizes each of the associated halfspaces. Passing to the cubical subdivision of C ensures that the action is without inversions.

Notation 2.4 (Translation length)**.** Let C be a CAT(0) cube complex. If $g \in \mathrm{Aut}(C)$ fixes a point, let $\|g\| = 0$. If $g \in \mathrm{Aut}(C)$ nontrivially stabilizes a combinatorial geodesic $\tilde{A}$ in C, then $\|g\|$ denotes the number of $\langle g \rangle$–orbits of 1–cubes in $\tilde{A}$. Note that if G acts on C without inversions, then $\|g\|$ is well-defined for all $g \in G$. Since we will consider actions on multiple cube complexes, we adopt the following convention on subscripts: when g is acting on $C_\heartsuit$, then we use the notation $\|g\|_\heartsuit$.

Observation 2.5. Suppose that $C_\heartsuit$ and $C_\diamondsuit$ are CAT(0) cube complexes, and G acts on $C_\heartsuit$ and $C_\diamondsuit$. Let $C = C_\heartsuit \times C_\diamondsuit$. Then, with respect to the product action $G \to \mathrm{Aut}(C)$,

$$\|g\|_C = \|g\|_\heartsuit + \|g\|_\diamondsuit$$

for all $g \in G$. Moreover, $\|g^d\|_\heartsuit = |d|\|g\|_\heartsuit$ for all $g \in G$ and $d \in \mathbb{Z}$.

2.3. **Cutting and chopping.** We now recall a special case of the notion of a *wallspace*, introduced in [27], and refer the reader to [28] for a more detailed discussion.

A *wallspace* $(\widetilde{X}, \mathcal{W})$ consists of a metric space $\widetilde{X}$ with a set $\mathcal{W}$ of connected subspaces $\widetilde{W}$, called *walls*, so that $\widetilde{X} - \widetilde{W}$ has exactly two components, called *halfspaces*, with the additional property that for all $x, y \in \widetilde{X}$, there are finitely many $\widetilde{W} \in \mathcal{W}$ such that x, y lie in different components of $\widetilde{X} - \widetilde{W}$ (in which case $\widetilde{W}$ *separates* x, y). An *automorphism* of $(\widetilde{X}, \mathcal{W})$ is an isometry $g : \widetilde{X} \to \widetilde{X}$ so that $g\widetilde{W} \in \mathcal{W}$ for $\widetilde{W} \in \mathcal{W}$; the group of automorphisms of $(\widetilde{X}, \mathcal{W})$ is denoted $\mathrm{Aut}(\widetilde{X}, \mathcal{W})$.

Definition 2.6 (Cutting, parallel)**.** Let $(\widetilde{X}, \mathcal{W})$ be a wallspace and let $g \in \mathrm{Aut}(\widetilde{X}, \mathcal{W})$. Let $A : \mathbb{R} \to \widetilde{X}$ be an embedding with g–invariant image. Then $\widetilde{W} \in \mathcal{W}$ *cuts* g if there exists $x \in \mathbb{R}$ so that $A((x, \infty))$ and $A((-\infty, x))$ lie in different halfspaces associated to $\widetilde{W}$ and $\{g^n\widetilde{W} \cap A : n \in \mathbb{Z}\} = A(y\mathbb{Z})$ for some $y \in \mathbb{R} - \{0\}$.

We call $g, h \in \mathrm{Aut}(\widetilde{X}, \mathcal{W})$ *parallel* if: each $W \in \mathcal{W}$ cuts g if and only if W cuts h.

The CAT(0) cube complex $C = C(\widetilde{X}, \mathcal{W})$ *dual to* the wallspace $(\widetilde{X}, \mathcal{W})$ was defined in [36]. Denoting by $\widehat{\mathcal{W}}$ the set of halfspaces, the 0–cubes of C are maps $c : \mathcal{W} \to \widehat{\mathcal{W}}$ sending each wall to one of the two associated halfspaces, subject to: $c(\widetilde{W}) \cap c(\widetilde{W}') \neq \emptyset$ for $\widetilde{W}, \widetilde{W}' \in \mathcal{W}$, and for all $x \in \widetilde{X}$, we have $|\{\widetilde{W} : x \notin c(\widetilde{W})\}| < \infty$. The 0–cubes c, c' are joined by a 1–cube if and only if there is a unique wall $\widetilde{W}$ with $c(\widetilde{W}) \neq c'(\widetilde{W})$, and higher-dimensional cubes are added when their 1–skeleta appear.

If G acts on $(\widetilde{X}, \mathcal{W})$, then there is an induced G–action on C defined as follows: for $g \in G$ and $c \in C^0$, we have $(gc)(\widetilde{W}) = g \cdot c(g^{-1}\widetilde{W})$ for all $\widetilde{W} \in \mathcal{W}$. This action clearly takes hyperplanes to hyperplanes, and there is a G–equivariant bijection from $\mathcal{W}$ [respectively $\widehat{\mathcal{W}}$] to the set of hyperplanes [respectively, halfspaces] in C.

We now state two lemmas, the first of which is [38, Lemma 2.1]. The second follows by considering hyperplanes intersecting axes in the dual cube complex.

Lemma 2.7 (Cut-wall criterion)**.** *Let G act on a wallspace $(\widetilde{X}, \mathcal{W})$ and suppose that each $g \in G - \{1_G\}$ is cut by some wall in $\mathcal{W}$. Then G acts freely on $C(\widetilde{X}, \mathcal{W})$.*

Lemma 2.8 (Cut-walls and translation length in $C(\widetilde{X}, \mathcal{W})$)**.** *Let G act on the wallspace $(\widetilde{X}, \widetilde{W})$ and let $g \in G$. Let $\|g\|_{\mathcal{W}}$ be the number of $\langle g \rangle$–orbits of walls that cut g. Suppose, moreover, that for all $\widetilde{W} \in \mathcal{W}$ such that $g\widetilde{W} = \widetilde{W}$, each halfspace associated to $\widetilde{W}$ is preserved by g. Then the translation length $\|g\|$ of g on C is defined and $\|g\| = \|g\|_{\mathcal{W}}$.*

Definition 2.9 (g–chopping wall)**.** Let $(\widetilde{X}, \mathcal{W})$ be a wallspace and let $g \in \mathrm{Aut}(\widetilde{X}, \mathcal{W})$. A wall $\widetilde{W} \in \mathcal{W}$ is *g–chopping* if $\widetilde{W}$ cuts g and there exists $n > 0$ such that the walls $g^{kn}\widetilde{W}, k \in \mathbb{Z}$ are pairwise disjoint.

A CAT(0) cube complex C is a wallspace where the walls are the hyperplanes. If G acts on C and $g \in G$, then the hyperplane $\widetilde{H} \subset C$ is *g–chopping* if $\widetilde{H}$ is a g–chopping wall in C. Moreover, if G acts on $(\widetilde{X}, \mathcal{W})$, then if $g \in G$ and $\widetilde{W} \in \mathcal{W}$ is a g–chopping wall, then the corresponding hyperplane of $C(\widetilde{X}, \mathcal{W})$ is g–chopping.

2.4. **Cubical small-cancellation theory.** We review the following background from [40], which plays a role in the proof of Theorem 5.2 below. Throughout, $\widetilde{X}$ denotes a CAT(0) cube complex and X a nonpositively-curved cube complex. For a hyperplane $\widetilde{U}$ of $\widetilde{X}$, we denote by $N(\widetilde{U})$ its *carrier*, i.e. the union of all closed cubes intersecting $\widetilde{U}$. We do the same for (immersed) hyperplanes in X. The *systole* $\|X\|$ is the infimal length of an essential combinatorial closed path in X.

2.4.1. *Cubical presentations and pieces.*

Definition 2.10 (Cubical presentation)**.** A *cubical presentation* $\langle X | \{Y_i : i \in I\}\rangle$ consists of non-positively curved cube complexes X and $\{Y_i\}$, and local isometries $\{Y_i \to X : i \in I\}$. We use the notation X^* for the cubical presentation above. As a topological space, X^* is X with a cone on each Y_i.

Definition 2.11 (Piece)**.** A *cone-piece* of X^* in Y_i is a component of $\widetilde{Y}_i \cap g\widetilde{Y}_j$, where $g \in \pi_1(X)$, excluding the case where $i = j$ and $g \in \mathrm{Stab}(\widetilde{Y}_i)$. A *wall-piece* of X^* in Y_i is a component of $\widetilde{Y}_i \cap N(\widetilde{U})$, where $\widetilde{U}$ is a hyperplane that is disjoint from $\widetilde{Y}_i$. A *piece-path* in Y is a path in a piece of Y.

2.4.2. *Simplified $B(6)$ condition.*

Definition 2.12 ($C'(\alpha)$ cubical presentation)**.** The cubical presentation $X^* = \langle X \mid \{Y_i : i \in I\}\rangle$ satisfies the *$C'(\alpha)$ condition* if $|P| \leq \alpha|S|$ for each geodesic piece-path P and each essential closed path $S \to Y_i$ with P a subpath of S.

The following is a simplified form of the $B(6)$ condition described in [40]; see Definition 5.1 and Remark 5.2.

Definition 2.13 (Simplified $B(6)$)**.** Let $X^* = \langle X \mid \{Y_i : i \in I\}\rangle$ be a cubical presentation satisfying the $C(\frac{1}{12})$ condition. Then X^* satisfies the *$B(6)$ condition* if the following hold:

(1) for each $i \in I$, there is an equivalence relation $\sim_i$ on the hyperplanes of Y_i.
(2) for each $i \in I$ and each pair of (possibly equal) hyperplanes U, V of Y_i with $U \sim_i V$, the hyperplanes U, V do not cross or osculate;

(3) for each $i \in I$, and each $\sim_i$–class $[U]$ of hyperplanes, $W = \cup_{V \in [U]} V$ is a *wall*: the space $Y_i - W$ consists of two subspaces $\overleftarrow{Y}_i, \overrightarrow{Y}_i$ such that $\mathrm{Cl}(\overleftarrow{Y}_i) \cap \mathrm{Cl}(\overrightarrow{Y}_i) = W$;
(4) if $P \to Y_i$ is a path that is the concatenation of at most 7 piece-paths and P starts and ends on the carrier $N(U)$ of a wall then P is path-homotopic into $N(U)$;
(5) $\mathrm{Aut}(Y_i \to X)$ preserves the wallspace structure.

The $B(6)$ condition provides a wallspace structure on $\widetilde{X^*}$ as follows: two hyperplanes of $\widetilde{X^*}$ are *elementary equivalent* if the intersect a lift $Y_i \hookrightarrow \widetilde{X^*}$ of some Y_i in equivalent hyperplanes. The transitive closure of this relation is an equivalence relation on the hyperplanes of $\widetilde{X^*}$ and the union of the hyperplanes in each class is a wall. The CAT(0) cube complex dual to this wallspace is the cube complex *associated* to X^*; observe that $\pi_1 X^*$ acts on this cube complex.

2.4.3. *Free action of $\pi_1 X^*$ on the associated cube complex.*

Definition 2.14 (Proximate)**.** Let $\langle X \mid \{Y_i : i \in I\}\rangle$ be a cubical presentation satisfying Definition 2.13.(2),(3). A hyperplane v in Y_i is *proximate* to a hyperplane u if there are 1-cells e_v, e_u dual to v, u such that either e_v, e_u have a common endpoint, or e_u has an endpoint on a piece in Y_i containing e_v. (Note that the first possibility includes the cases where v, u are equal, cross, or osculate.) A wall V in Y_i is *proximate* to a hyperplane u if some hyperplane v of V is proximate to u.

Theorem 2.15. *If $\langle X \mid \{Y_i\}_i\rangle$ satisfies the following properties then $\pi_1 X^*$ acts with torsion stabilizers on its associated dual cube complex.*

(1) *X^* satisfies the $B(6)$ condition.*
(2) *The following holds for each $Y \in \{Y_i\}$. Let κ be a geodesic of Y whose first and last edges are dual to hyperplanes H_p and H_q of Y. Suppose there are distinct hyperplanes w_1, w_1' in the same wall Y, where κ traverses an edge dual to w_1 and that w_1' is proximate to H_q. Then there is a wall w_2 in Y that separates H_p, H_q but is not proximate to H_p or H_q.*

We will apply Theorem 2.15 in the case where each Y_i is a specific CAT(0) cube complex mapping by a local isometry to X. We also Theorem 5.20 in [40], i.e.:

Theorem 2.16. *Let X^* be a $B(6)$ presentation and let W be a wall in $\widetilde{X}^*$. If H_1, H_2 are hyperplanes in W, and Y is a cone, then $H_1 \cap Y$ and $H_2 \cap Y$ lie in the same wall of Y.*

2.5. **Graphs of groups with cyclic edge groups.** The following easy lemma about groups acting on trees is used in the proof of Theorem 5.2, to show that the automorphism group of a relator in a certain cubical presentation preserves the chosen wallspace structure, which is one of the requirements of the $B(6)$ condition.

Lemma 2.17. *Let H act without inversions on a tree $\mathcal{T}$. Suppose that for each edge e of $\mathcal{T}$, the following hold:*

(1) *the edge group $H_e = \langle h_e\rangle$ is a maximal $\mathbb{Z}$ subgroup;*
(2) *if $kh_e^p k^{-1} = h_e^q$ for some $k \in H, p, q \neq 0$ then $kh_e = h_e k$.*

Let L be a combinatorial line in T with stabilizer $K \leq H$. Then both of the following hold:

- *every element of K either acts as a translation on L or fixes L pointwise;*
- *one of the following holds:*
  - *$K \cong \{1\}$;*
  - *$K \cong \mathbb{Z}$ and K is generated by a translation of L;*

– *each $k \in K$ centralizes h_e for each edge e of L.*

Proof. Denote by $\mathrm{Aut}(L) \cong \mathbb{Z}_2 * \mathbb{Z}_2$ the group of combinatorial automorphisms of L not inverting edges. Suppose that v is a vertex of L with incident edges e, e'. Let $H_e = \langle h_e \rangle$ and $H_{e'} = \langle h_{e'} \rangle$. Suppose $x \in K$ exchanges e, e'. Then $x^2 \in \langle h_e \rangle$, but x is not, contradicting maximality of the edge groups. Hence $\mathrm{image}(K \to \mathrm{Aut}(L))$ is trivial or infinite cyclic, which proves the first assertion.

Note that $\ker(K \to \mathrm{Aut}(L)) = \cap_{e \in \mathrm{Edges}(L)} H_e$. Suppose that $k \in K$ and $h \in \ker(K \to \mathrm{Aut}(L)) - \{1\}$. Then for each edge e of L, there exists $n \neq 0$ so that $h = h_e^n$. Since k is a translation, $[k, h]$ stabilizes e, i.e. there exists $\ell \in \mathbb{Z}$ so that $[k, h_e^n] = h_e^\ell$, i.e. $k h_e^n k^{-1} = h_e^{\ell+n}$. Hence $[k, h_e] = 1$ by hypothesis (2). Hence either k centralizes each h_e or $\ker(K \to \mathrm{Aut}(L)) = \{1\}$. □

3. Proof of Theorem A

We now prove Theorem A, using Theorems 4.1 and 5.2; the rest of the paper is devoted to proving these statements.

Proof of Theorem A. Let $G = F \rtimes_\Phi \langle t \rangle$ and let Z_t denote the centralizer of t in G. Proposition 2.3 provides a finitely-generated free supergroup F' of F and a sequence $Z_t = G'_0 \subset G'_1 \subset \cdots \subset G'_n = F' \rtimes_\Phi \mathbb{Z}$ of F'_i-by-$\langle t \rangle$ groups defined as follows. For each $i \geq 0$, the group G'_{i+1} is the following cyclic HNN extension of G'_i:

$$G'_{i+1} \quad = \quad \langle G'_i, e_{i+1} \mid e_{i+1} t e_{i+1}^{-1} = p_{i+1} t \rangle \tag{1}$$

where e_{i+1} corresponds to the new edge of a polynomial stratum and p_{i+1} is an element of $F_i \subset G'_i$. Note that it is possible that $p_k = p_\ell$ for $k \neq \ell$. It suffices to produce a free action of G'_n on a CAT(0) cube complex, since G is virtually contained in G'_n. Let $J_i = \{k \,:\, y_k p_k y_k^{-1} \in G'_i \text{ for some } y_k \in G'_n\}$.

For each $i \geq 0$ and each $j \in J_{i+1} - J_i$, let $\widetilde{A}_j$ be the axis of $y_j(p_j t)y_j^{-1} \in G'_{i+1}$ in the Bass-Serre tree of the splitting of G'_{i+1} in Equation (1). For each $y_j(p_j t)y_j^{-1}$, we have a vertex v of $\widetilde{A}_j$ whose incident edge-groups in $\widetilde{A}_j$ are subgroups $\langle r_j \rangle, \langle s_j \rangle$. The pair (r_j, s_j) is the *diverging pair* associated to $y_j(p_j t)y_j^{-1}$. (Since we will arrange for these elements to be non-parallel in the eventual cubulation.)

For each i, let M_i be a finite set of pairs (r_j, s_j) which is finite up to simultaneous conjugacy in G'_{i+1} and has the following properties:

(1) M_i contains each diverging pair (r_j, s_j) for which there exists g such that $r_j^g, s_j^g \in G'_i$;
(2) if $(r, s) \in M_{i+1}$ and r, s stabilize vertices v_r, v_s of the Bass-Serre tree, and $\mathsf{d}(v_r, v_s) \geq 2$, then there are adjacent edges e, e' separating v_r, v_s and $(r', s') \in M_i$ so that r' stabilizes e and s' stabilizes e';
(3) if $(r, s) \in M_{i+1}$ and r, s stabilize vertices v_r, v_s with $\mathsf{d}(v_r, v_s) = 1$, then there exists $r' \in \mathrm{Stab}(v_r) \cap \mathrm{Stab}(v_s)$, so that $(r, r') \in M_i$.

We now prove the following claim by induction on i.

Claim 1. *G'_i acts freely on a CAT(0) cube complex C_i, and this action satisfies:*

- *$\|t\|_i = \|p_j t\|_i$ for all $j \in J_i$.*
- *There is a t-chopping wall in C_i.*
- *Any $(r_j, s_j) \in M_i$ is non-parallel in C_i.*

The base case, $i = 0$, holds by Theorem 4.1, since $G'_0 = Z_t$. Assume that the claim holds for some $i \geq 0$. By Lemma 5.14, the subgroup $\langle t \rangle$ is separable in G'_{i+1}. Note that G'_{i+1} satisfies the

remaining hypotheses of Theorem 5.2, where L_0 is M_i modulo conjugacy and L_1 is $M_{i+1} - M_i$ modulo conjugacy; in particular, apply Lemma 5.15. The sets $\{p_j : j \in J\}, \{q_k : k \in K\}$ from Theorem 5.2 are respectively $\{p_j : j \in J_i\}$ and $\{p_j : j \in J_{i+1} - J_i\}$. Theorem 5.2 now yields the claim for G'_{i+1}. □

4. Balanced cubulation of $F \times \mathbb{Z}$

The goal of this section is to prove Theorem 4.1, which includes Theorem B:

Theorem 4.1 (Balanced cubulation of $F \times \mathbb{Z}$)**.** *Let $\mathcal{P}$ be a finite set of nontrivial elements of the finitely-generated free group F. Then there exists a free action of $G = F \times \langle t \rangle$ on a CAT(0) cube complex C so that $\|pt\| = \|t\|$ for all $p \in \mathcal{P}$. Moreover, for each $p \in \mathcal{P}$ and any finite set $\{g_1, \ldots, g_\ell\} \subset F$:*

- *t and $(pt)^h$ are non-parallel for all $h \in F$;*
- *pt and $(pt)^{g_l}$ are non-parallel unless they are equal, for $1 \leq l \leq \ell$.*

Proof. Lemma 4.3 provides a (non-free) action of G on a CAT(0) cube complex $C_\bullet$ such that $\|pt\|_\bullet = \|t\|_\bullet > 0$ for each $p \in \mathcal{P}$. Lemma 4.2 provides a (non-free) action of F on a CAT(0) cube complex $C_\copyright$ such that $\|p\|_\copyright = 0$ for all $p \in \mathcal{P}$ and such that any $f \in F$ either satisfies $\|f\|_\copyright > 0$ or f is conjugate into $\langle p \rangle$ for some $p \in \mathcal{P}$. Extend this to an action of $G = F \times \langle t \rangle$ with t acting trivially. Let $C = C_\bullet \times C_\copyright$ and let G act using the product action. Then for each $p \in \mathcal{P}$,

$$\begin{aligned} \|pt\| &= \|pt\|_\bullet + \|pt\|_\copyright \\ &= \|t\|_\bullet + \|p\|_\copyright = \|t\|. \end{aligned}$$

To see that this action is free, let $f \in F$ and $d \in \mathbb{Z}$, so that $\|ft^d\| = \|f\|_\copyright + |d|\|t\|_\bullet$. If $d \neq 0$, then $\|ft\| > 0$ since $\|t\|_\bullet > 0$. If f is not conjugate into $\langle p \rangle$ for any $p \in \mathcal{P}$, then $\|ft^d\| \geq \|f\|_\copyright > 0$ by Lemma 4.2. Otherwise, $d = 0$ and f is not conjugate into an element of $\mathcal{P}$, whence $\|ft\| = \|f\|_\bullet > 0$. Thus the action of G on C is free.

The moreover clause is handled by the corresponding clause in Lemma 4.3. □

Lemma 4.2 (Actions of F with specified $\mathbb{Z}$ stabilizers)**.** *Let $P_1, \ldots, P_r$ be distinct maximal cyclic subgroups of the free group F. Then there exists an action of F on a CAT(0) cube complex $C_\copyright$ such that:*

1. *F acts with no inversions in hyperplanes of $C_\copyright$;*
2. *each P_i fixes a 0–cube;*
3. *for each $c \in C^0_\copyright$, either $\mathrm{Stab}_F(c) = 1_F$ or $\mathrm{Stab}_F(c)$ is conjugate to some P_i.*

In particular, each $f \in F$ either acts hyperbolically on $C_\copyright$ or is conjugate into some element of $\mathcal{P}$.

Proof. Let $f \in F$ be *simple* if there does not exist $g \in F$ and $n > 1$ such that $f = g^n$. Let $\mathcal{S} \subset F$ be a set of cyclically reduced simple elements such that each maximal cyclic subgroup of F is conjugate to $\langle f \rangle$ for some unique $f \in \mathcal{S}$. Assume that $\mathcal{S}$ contains a generator of each P_i.

Let $f : \mathbb{N} \to \mathcal{S}$ be an enumeration of $\mathcal{S}$ such that $f(i)$ is a generator of P_i for $1 \leq i \leq r$. Given $g \in F$, let g' be the unique element of $\mathcal{S}$ generating the maximal cyclic subgroup containing g and let $\#(g) = f^{-1}(g')$.

Below, we shall choose a sequence (q_n) of natural numbers such that $q_n > \max_{1 \leq i \leq n}\{12, 6|f(i)|\}$.

For $n \geq 1$, let

$$F_n = F/\langle\langle f(1)^{q_n}, \ldots, f(n)^{q_n} \rangle\rangle.$$

The hypothesized lower bound on q_n implies that F_n admits a $C'(\frac{1}{6})$ small-cancellation presentation for all $n \in \mathbb{N}$. Hence each F_n acts properly and cocompactly on a CAT(0) cube complex C_n, as shown in [39].

Moreover, by choosing (q_n) to grow sufficiently quickly, we can ensure that $f(q_{\#(g)-1})$ has infinite order in $F_{\#(g)-1}$ for all $g \in F$. Hence for each n the element $f(n)$ acts hyperbolically on C_{n-1} and elliptically on C_m for $m \geq n$.

By passing to the cubical subdivision, if necessary, we may assume that each $g \in F$ has a well-defined translation length on each C_n (in the sense of Notation 2.4), and that each elliptic element fixes a 0–cube.

Let $C_{\copyright} \subset \prod_{n \geq r} C_n$ be the *restricted infinite product*, i.e. the CAT(0) subcomplex which is the convex hull of the set of 0–cubes that lie on a finite path from a fixed basepoint. The product action of F on $\prod_{n \geq r} C_n$ (where F acts as F_n on C_n) preserves $C_{\copyright}$ since $\|g\|_{\copyright} = \sum_{n \geq r} \|g\|_n < \infty$, since $\|g\|_n = 0$ for $n \geq \#(g)$. Note that $\|g\|_{\copyright} > 0$ if and only if $\#(g) \geq r$, so each P_i is the stabilizer of a 0–cube and each 0–cube stabilizer is conjugate to some P_i. □

Lemma 4.3 (Non-free balancing of t, pt)**.** *Let $\mathcal{P}$ be a finite set of nontrivial elements of the free group F. Then there exists an action of $G = F \times \langle t \rangle$ on a CAT(0) cube complex $C_\bullet$ such that:*

1. $\|p\|_\bullet > 0$ *for all* $p \in \mathcal{P}$*;*
2. $\|pt\|_\bullet = \|t\|_\bullet$ *for all* $p \in \mathcal{P}$*.*

Moreover, for each $p \in \mathcal{P}$ and any finite set $\{g_1, \ldots, g_\ell\} \subset F$:

- *t and $(t^n p)^h$ are non-parallel for all $h \in F$ and $n \in \mathbb{Z}$;*
- *pt and $(pt)^{g_l}$ are non-parallel unless they are equal, for $1 \leq l \leq \ell$.*

We emphasize that the lemma does **not** produce a free action on a cube complex.

Proof of Lemma 4.3. If the conclusion holds with $\mathcal{P}$ replaced by a finite superset, then it holds for $\mathcal{P}$, so we may assume, conjugating if necessary, that $\mathcal{P}$ is partitioned into finite sets $\mathcal{P}_1, \ldots, \mathcal{P}_r$ such that for each i, there exists $p_i \in \mathcal{P}_i$ so that

$$\mathcal{P}_i = \left\{p_i^{\pm 1}, p_i^{\pm 2}, \ldots, p_i^{\pm m}\right\}$$

for some fixed $m > 0$, and if $p_i^{n_i} = p_j^{n_j}$ for some $n_i, n_j \neq 0$, then $i = j$ and $n_i = n_j$.

Let B be a bouquet of circles with $\pi_1 B \cong F$ and let S be a bouquet of one circle based at the vertex with $\pi_1 S \cong \langle t \rangle$. For each i, we also let $p_i \in F$ denote the immersed closed combinatorial path representing that element. Let $\widehat{B} \to B$ be a finite regular cover in which each p_i^m embeds. Indeed, we may choose m ex post facto. For each i, let $T_i \subset \widehat{B} \times S$ be the torus $p_i^m \times t$, where $t \to S$ is the closed combinatorial path at the vertex of S representing the element t in $F \times \langle t \rangle \cong \pi_1 B \times \pi_1 S$. We will use the $\{\pm 1\} \times \{\pm 1\}$ action on T_i generated by combinatorial reflections of the factors at their basepoints.

Walls in $\widehat{B} \times S$: Let $s > m$ be an integer. For each $e \in \text{Edges}(\widehat{B})$ and each $z \in \{\pm 1\} \times \{\pm 1\}$, we construct an immersed wall $W_{e,z} \to \widehat{B} \times S$ as follows. Define geodesic arcs $\alpha_0, \ldots, \alpha_s$ in $[0,1] \times [0,1]$ as follows: let α_0 join $\left(\frac{1}{2}, 0\right)$ to $\left(1, \frac{1}{2s}\right)$, let α_k join $\left(0, \frac{2k-1}{2s}\right)$ to $\left(1, \frac{2k+1}{2s}\right)$ for $1 \leq k \leq s-1$, and let α_s join $\left(0, \frac{2s-1}{2s}\right)$ to $\left(\frac{1}{2}, 1\right)$. Let $R_{e,s}$ be the image of $\cup_{k=0}^{s} \alpha_k$ under $[0,1] \times [0,1] \to e \times S$.

The wall $W_{e,z}$ is formed from $zR_{e,s}$ so that $W_{e,z}$ is connected and satisfies:

- $W_{e,z} \cap (e \times S) = zR_{e,s}$;
- $W_{e,z}$ intersects each 2–cell of $(\widehat{B} - \text{Int}(e)) \times S$ (with the product cell-structure) in a set of s evenly spaced geodesic arcs parallel to an edge of $\widehat{B}$. See Figure 2.

Observe that $W_{e,z} \subset B \times S$ is *two-sided* since there is an embedding $W_{e,z} \times [-\epsilon, \epsilon] \to \widehat{B} \times S$ such that the restriction to $W_{e,z} \times \{0\}$ is the original inclusion $W_{e,z} \hookrightarrow \widehat{B} \times S$. Finally, observe that $W_{e,z}$ is transverse to the 1–skeleton of $\widehat{B} \times S$.

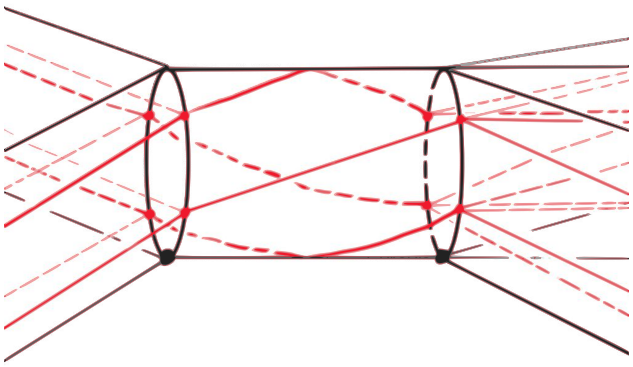

FIGURE 2. An immersed wall $W_{e,z}$ in $\widehat{B} \times S$, with $s = 4$. The cylinder in the center is $e \times t$, and its horizontal edge is parallel to e.

Counting intersections: We now count points of intersection between the immersed walls $W_{e,z} \hookrightarrow \widehat{B} \times S \to B \times S$ and the elements $t, p_i^j t$.

Elements of $\mathcal{P}$: For $1 \leq i \leq r$, the embedded path p_i^m intersects the wall $W_{e,z}$ transversely in a single point – the midpoint of e – if and only if p_i^m traverses e. Otherwise, $p_i^m \cap W_{e,z} = \emptyset$. Hence, since each edge yields 4 walls, the total number of intersection points of the path p_i^m with the various walls is $4|p_i^m|$.

Lifting walls to count intersections in finite covers: For each $d \geq 1$, let $\widehat{S}_d \to S$ be the connected d–fold cover, and consider the cover $\widehat{B} \times \widehat{S}_d \to \widehat{B} \times S \to B \times S$. For each edge e of $\widehat{B}$, and each $z \in \{\pm 1\} \times \{\pm 1\}$, let $\widehat{W}_{e,z}$ be the preimage of $W_{e,z}$ under the projection $\widehat{B} \times \widehat{S}_d \to \widehat{B} \times S$. Note that $\widehat{W}_{e,z}$ is an embedded, connected 2–sided immersed wall consisting of ds parallel copies of $\widehat{B} - \mathrm{Int}(e)$, together with ds arcs of slope $\frac{1}{s}$, all lying in the cylinder $e \times t^d$.

Powers of t: For any $d \geq 1$, the embedded cycle $t^d \in \widehat{B} \times \widehat{S}_d$ intersects each $\widehat{W}_{e,z}$ in exactly ds points, for a total of $4ds|\mathrm{Edges}(\widehat{B})|$ intersection points between t^d and all walls $\{\widehat{W}_{e,z}\}$ in $\widehat{B} \times \widehat{S}_d$.

Powers of $p_i t$: Let $i \in \{1, \ldots, r\}$ and let $n \in \{\pm 1 \ldots, \pm m\}$. Let $d = \frac{m}{gcd(m,n)}$ and let $d' = \frac{n}{gcd(m,n)}$ so that $dn = d'm$. Let e be an edge of $\widehat{B}$. Working in the cover $\widehat{B} \times \widehat{S}_d$, we will compute the minimal number of intersections between each $\widehat{W}_{e,z}$ and a closed path homotopic to $(p_i^n t)^d$.

First, consider the case in which $(p_i^n t)^d$ does not traverse e. In this case, $(p_i^n t)^d$ is homotopic to a path $\widehat{p_i^n} t^d$ that travels along t^d in the cylinder $e \times t^d$ but travels within $(\widehat{B} - \mathrm{Int}(e)) \times \{0\}$ outside of that cylinder. Thus, up to homotopy, $(p_i^n t)^d$ intersects $\widehat{W}_{e,z}$ in ds points. So, fixing e and letting z vary in $\{\pm 1\} \times \{\pm 1\}$, we see $4ds$ such intersections.

Second, consider the case in which $(p_i^n t)^d$ traverses e. Since $nd = md'$, the path $(p_i^n t)^d$ in the torus $p_i^m \times \widehat{S} \subset \widehat{B} \times \widehat{S}$ is homotopic to a closed geodesic γ of slope $\frac{d}{d'}$, while $\widehat{W}_{e,(i,j)}$ intersects that torus in a closed path homotopic to a closed geodesic of slope $\frac{i}{js}$ for $i, j \in \{\pm 1\}$. Hence, relative to the basis corresponding to the product structure, these correspond to vectors $\begin{pmatrix} d' \\ d \end{pmatrix}$ and $\begin{pmatrix} is \\ j \end{pmatrix}$.

Hence

$$|\gamma \cap \cup_z \widehat{W}_{e,z}| = \sum_{i,j\in\{\pm 1\}} \left|\det \begin{pmatrix} is & d' \\ j & d \end{pmatrix}\right| = 4ds.$$

The second equality holds since otherwise $sd \leq |d'|$, whence $md < |d'|$ since $m < s$. But then $|n|d \leq md < |d'|$, which contradicts that $|n|d = m|d'| \geq |d'|$.

We have thus shown that the total number of intersections of any $(p_i^n t)^d$ with $\cup_{e,z}\widehat{W}_{e,z}$ is, up to homotopy, $4ds|\mathrm{Edges}(\widehat{B})|$.

Conclusion: Observe that the set of immersed walls $W_{e,z}$ is invariant under the action of $G \cong \pi_1 B \times \pi_1 S$ on $\widehat{B} \times S$. Hence the universal cover $\widetilde{B} \times \widetilde{S}$ contains a G–invariant set $\mathcal{W}$ of walls consisting of components of preimages of the various $W_{e,z}$ under the projection $\widetilde{B} \times \widetilde{S} \to \widehat{B} \times S$. Since each $W_{e,z}$ is 2–sided, G acts without inversions in hyperplanes on the cube complex $C_\bullet$ dual to the wallspace $(\widetilde{B} \times \widetilde{S}, \mathcal{W})$.

The preceding discussion and Lemma 2.8 thus show that, for all $p \in \mathcal{P}$ and $d \geq 1$ chosen as above in terms of p,

$$\|t^d\|_\bullet = \|(pt)^d\| = 4ds|\mathrm{Edges}(\widehat{B})|,$$

so that $\|t\|_\bullet = \|pt\|_\bullet$. Similarly, $\|p\|_\bullet > 0$ for all $p \in \mathcal{P}$.

Non-parallelism of distinct conjugates: Non-parallelism of t^g and $(pt)^h$ follows by considering any $W_{e,z}$ where p does not traverse e.

Now consider distinct conjugates $pt, (pt)^{g_l}$. Without loss of generality, we may assume that $\widehat{B}$ is a regular cover with the property that the g_l–translate of the elevation of p is distinct from p, when $[g_l, p] \neq 1$. Let e be an edge traversed by exactly one of the paths p or its g_l–translate. Then any wall $W_{e,z}$ distinguishes the parallelism types of $pt, (pt)^{g_l}$. □

Remark 4.4 (Metrically proper action)**.** The construction of C in Theorem 4.1 can be slightly modified (by cubical subdivisions in the $C_{\copyright}$ factor) so that the action of $F \times \langle t \rangle$ on C is metrically proper. In fact we now show that if $(f_j t^{k_j})$ is a sequence of distinct elements in $F \times \langle t \rangle$, then $\|f_j t^{k_j}\|_C \to \infty$ as $j \to \infty$.

First, by subdividing C_n in the proof of Lemma 4.2 so that each hyperplane is replaced by n parallel copies of itself, one can arrange for the action of F on $C_{\copyright}$ to have the following "weak metric properness" property: if (f_j) is a sequence in F and $\|f_j\|_{\copyright}$ is bounded, then f_j eventually lies in some conjugate of some P_i.

Indeed, for each j, let $b_j = \#(f_j) - 1$, so that f_j acts hyperbolically on C_{b_j}. Thus $\|f_j\|_{\copyright} \to \infty$ since the translation length of f_j on C_{b_j}, and hence on $C_{\copyright}$, is at least b_j because of the subdivision. Thus $\|f_j\|_{\copyright}$ is bounded only if, after discarding finitely many j, each f_j maps to a torsion element of F_r, i.e. is conjugate into some P_i.

Second, consider the action of $F \times \langle t \rangle$ from Theorem 4.1 and let $(f_j t^{k_j})$ be a sequence in $F \times \langle t \rangle$. Then for each j, we have $\|f_j t^{k_j}\| = \|f_j\|_{\copyright} + \|f_j t^{k_j}\|_\bullet$. Hence, if $\lim_j \|f_j t^{k_j}\| \neq \infty$, then by passing to a subsequence we can assume that $\|f_j\|_{\copyright}$ is bounded and thus each f_j is conjugate into some fixed P_i. Since we are focusing on translation lengths, by conjugation it suffices to consider sequences of the form $(p^{\ell_j} t^{k_j})$ where $p \in P_i$ and $\ell_j, k_j \in \mathbb{Z}$.

We now show that $\|p^{\ell_j} t^{k_j}\|_\bullet \to \infty$. Since $\|pt\|_\bullet = \|t\|_\bullet$, we have:

$$\|p^{\ell_j} t^{k_j}\|_\bullet = \|(pt)^{\ell_j} t^{k_j - \ell_j}\|_\bullet \geq \big|\|(pt)^{\ell_j}\|_\bullet - \|t^{\ell_j - k_j}\|_\bullet\big| = \big||\ell_j| - |\ell_j - k_j|\big|\|t\|_\bullet$$

If $\big||\ell_j| - |\ell_j - k_j|\big| \to \infty$, then this gives $\|p^{\ell_j} t^{k_j}\| \to \infty$ as required. Otherwise, passing to a subsequence, $|k_j|$ is bounded by some K and we instead use the estimate $\|p^{\ell_j} t^{k_j}\| \geq \big||\ell_j|\|p\|_\bullet - K\|t\|_\bullet\big|$ and the fact that $\|p\|_\bullet > 0$ by Lemma 4.3.

5. Cubulating balanced cyclic HNN extensions

The goal of this section is to prove Theorem 5.2. In this section, given a group G and $t \in G$, we denote by $Z_G(t)$ the centralizer of t. Given a group G acting on a wallspace with set $\mathcal{W}$ of walls, we say that $g, h \in G$ are *$\mathcal{W}$–parallel* if the set of walls cutting g is equal to the set of walls cutting h.

Definition 5.1 (Exit property)**.** Let Γ be a graph of spaces, with fundamental group G', where each vertex space is a nonpositively-curved cube complex. Then the hyperbolic element $h \in G'$ has the *exit property* if the following holds. Let $\widetilde{\sigma}$ be the axis of h in the Bass-Serre tree. Then there exists a vertex $v \in \widetilde{\sigma}$ so that the corresponding vertex space contains a hyperplane that intersects exactly one of the edge spaces of $\widetilde{\sigma}$ incident to v. See Figure 3.

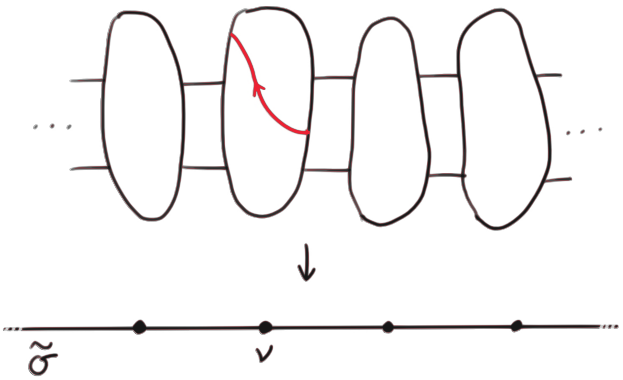

Figure 3. The exit property: the arrowed arc represents a hyperplane intersecting exactly one of the edge spaces of $\tilde{\sigma}$ incident to v.

Theorem 5.2. *Let G act freely on a CAT(0) cube complex C_0, let $t \in G - \{1\}$, and let $\{p_j\}_{j\in J}$ be a finite subset of G, with $p_1 \in \{p_j\}_{j\in J}$, i.e $1 \in J$. Let $G' = \langle G, e \mid t^e = p_1 t\rangle$, and let $\{q_k\}_{k\in K}$ be a finite subset of $G' - (G \cup Z_{G'}(t))$. Suppose any $g \in G'$ that commensurates t centralizes t. Fix a finite set $\{(r_\ell, s_\ell)\}_{\ell\in L_0}$ of pairs of nontrivial elliptic elements of G' and a finite set $\{(r_\ell, s_\ell)\}_{\ell\in L_1}$ such that r_ℓ, s_ℓ do not stabilize the same vertex of the Bass-Serre tree $\mathcal{T}$ of the splitting of G'. Suppose:*

1. *$\langle t\rangle$ is separable in G;*
2. *$\|p_j t\|_0 = \|t\|_0$ for each $j \in J$;*
3. *for all $\ell \in L_0$, the elements r_ℓ, s_ℓ stabilize a common vertex of $\mathcal{T}$ and are non-parallel in the corresponding translate of C_0;*
4. *for each $k \in K$, there exists $\ell \in L_0$ so that r_ℓ, s_ℓ stabilize adjacent edges e_{r_ℓ}, e_{s_ℓ} on the axis for $q_k t$;*
5. *for each $\ell \in L_1$, if r_ℓ, s_ℓ are both elliptic, then one of the following holds:*

- *if $\ell \in L_1$ and r_ℓ, s_ℓ stabilize vertices v_ℓ, w_ℓ of $\mathcal{T}$, and $\mathsf{d}_{\mathcal{T}}(v_\ell, w_\ell) \geq 2$, then there are adjacent edges e, e' separating v_ℓ, w_ℓ and $\ell' \in L_0$ so that $r_{\ell'}$ stabilizes e and $s_{\ell'}$ stabilizes e';*
- *if $\ell \in L_1$ and r_ℓ, s_ℓ stabilize vertices v_ℓ, w_ℓ with $\mathsf{d}_{\mathcal{T}}(v_\ell, w_\ell) = 1$, then there exists $\ell' \in L_0$ so that $r_{\ell'} \in \mathrm{Stab}(v_r) \cap \mathrm{Stab}(v_s)$ and $s_{\ell'} = r_\ell$.*

(6) *for each $\ell \in L_1$ with r_ℓ hyperbolic, r_ℓ is not conjugate into $Z_{G'}(t)$;*
(7) *C_0 has a t-chopping hyperplane.*

Then:

- *G' acts freely on a CAT(0) cube complex C_1.*
- *$\|t\|_1 = \|p_i t\|_1$ for each $i \in J$ and each $i \in K$.*
- *C_1 has a t-chopping hyperplane.*
- *r_ℓ, s_ℓ are not parallel in C_1 for all $\ell \in L_0 \cup L_1$.*

Remark 5.3 (Applying Theorem 5.2 in the free-by-$\mathbb{Z}$ case)**.** Theorem 5.2 is applied to a free-by-$\mathbb{Z}$ group G' in the proof of Theorem A. The splitting of G' as a graph of groups with cyclic edge groups and vertex groups acting on cube complexes is described in the proof of Theorem A, so to apply Theorem 5.2, one must check that the hypotheses are satisfied by mapping tori of polynomially growing free group automorphisms, which we do in Section 5.3.

Proof of Theorem 5.2. The proof is divided into several parts. First, by passing if necessary to the cubical subdivision of C_0, we can assume that G acts on C_0 without inversions, so that each nontrivial element stabilizes a bi-infinite combinatorial geodesic. Elements of G' are either *elliptic* or *hyperbolic* according to whether they fix a vertex of the underlying tree. The proof is summarized in Figure 4.

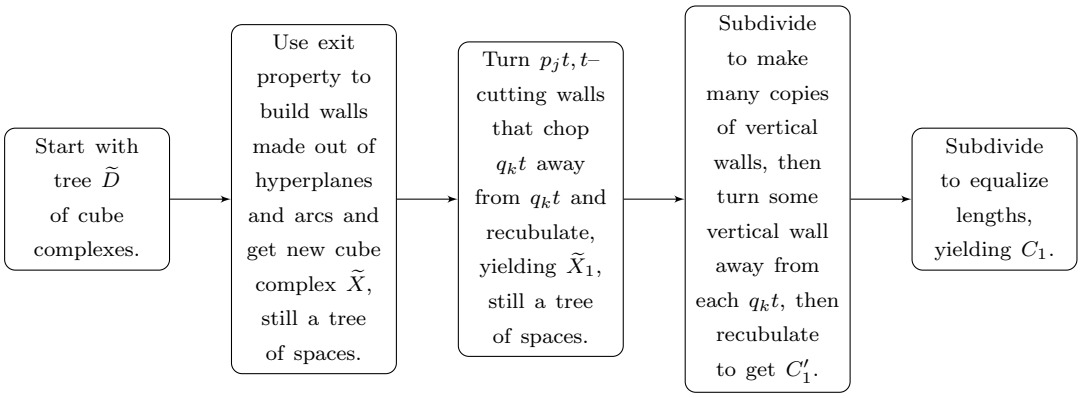
Start with tree $\widetilde{D}$ of cube complexes.
Use exit property to build walls made out of hyperplanes and arcs and get new cube complex $\widetilde{X}$, still a tree of spaces.
Turn $p_j t, t$–cutting walls that chop $q_k t$ away from $q_k t$ and recubulate, yielding $\widetilde{X}_1$, still a tree of spaces.
Subdivide to make many copies of vertical walls, then turn some vertical wall away from each $q_k t$, then recubulate to get C_1'.
Subdivide to equalize lengths, yielding C_1.

FIGURE 4. Plan of the proof of Theorem 5.2

The tree $\widetilde{D}$ of spaces: Let $\widetilde{A}$ be a t-invariant combinatorial geodesic in the cubical subdivision of $\widetilde{C}_0$. Choose $\widetilde{A}_1$ analogously for $p_1 t$. Let $A = \langle t \rangle \backslash \widetilde{A}$ and let $A_1 = \langle p_1 t \rangle \backslash \widetilde{A}_1$. Let $\overline{C}_0 = G \backslash C_0$. Let D be the quotient of $\overline{C}_0 \sqcup B \times [-1, 1]$, where B is a circle, formed by identifying each point of $B \times \{-1\}$ with the corresponding point of $\overline{C}_0$ using the map $B \cong A \to C_0$, and by identifying each point of $B \times \{1\}$ with the corresponding point in $\overline{C}_0$ using $B \cong A_1 \to \overline{C}_0$. Let $\widetilde{D} \to D$ be the universal cover, which is a tree of spaces with underlying tree $\mathcal{T}$ where each vertex space is a $G' \cong \pi_1 D$–translate of $\widetilde{C}_0$ and each edge space is a copy of the line $\widetilde{B}$. The preimage of an edge

of $\mathcal{T}$ under the map $\widetilde{D} \to \mathcal{T}$ is a strip $\widetilde{B} \times [-1,1]$ with $\widetilde{B} \times \{\pm 1\}$ identified with the appropriate copies of $\widetilde{A}, \widetilde{A}_1$ in the relevant vertex spaces.

A convenient finite cover: By Corollary 5.10, there exists a finite regular cover $\widehat{D} \to D$ splitting as a finite graph Γ of spaces whose vertex spaces are covers of C_0 and whose edge spaces are covers of B attached along elevations of A, A_1. The main property of $\widehat{D}$ is the following: let h_k generate $\langle q_k t \rangle \cap \pi_1 \widehat{D}$ and let $\tilde{\gamma}_k$ be an axis in $\widetilde{D}$ for $q_k t$. Then the path $\gamma_k = \tilde{\gamma}_k / \langle h_k \rangle$ projects to an embedded cycle σ_k in Γ.

Verifying the exit property: Each $q_k t$ has the exit property by (3),(4).

Horizontal and vertical walls: By equivariantly replacing each vertex space with an appropriate cubical subdivision and applying Corollary 5.13, using the above conclusion that each $q_k t$ has the exit property, we have a finite family $\mathcal{H}$ of *horizontal* immersed walls $W \to \widehat{D} \to D$ with the following properties:

(1) Each $W \in \mathcal{H}$ decomposes as a finite graph of spaces, whose vertex spaces are immersed hyperplanes in vertex spaces of $\widehat{D}$ and whose edge spaces are properly embedded arcs in edge spaces $\widehat{B} \times [-1,1]$ whose endpoints are intersections between immersed hyperplanes of the incident vertex spaces and the attaching circles $\widehat{B} \times \{\pm 1\} \to \widehat{D}$;

(2) For each hyperplane H of each vertex-space of $\widehat{D}$, there exists $W \in \mathcal{H}$ with $H \subset \widetilde{W}$ for some lift $\widetilde{W}$ of W.

(3) If $U \in \mathcal{H}$ has a lift $\widetilde{U}$ cutting gtg^{-1} [resp. $g(p_j t)g^{-1}$] and $q_k t$ for some $k \in K$ and $g \in G'$ [and $j \in J$], then $\widetilde{U}$ chops $q_k t$. Indeed, walls that cut $q_k t$ but fail to chop $q_k t$ would fail to avoid the $q_k t$–cycle in $\widehat{D}$, contradicting Corollary 5.13.

(4) For each $\ell \in L_1$ with r_ℓ, s_ℓ elliptic, stabilizing vertices $v_\ell, w_\ell \in \mathcal{T}$ respectively, there exists $W_\ell \in \mathcal{H}$ so that a lift $\widetilde{W}_\ell$ cuts r_ℓ but not s_ℓ or vice versa. We now explain how to arrange for this by replacing $\mathcal{H}$ with a larger set of walls. Note that adding walls to $\mathcal{H}$ does not affect the above listed properties.

First, let L_1' be the set of $\ell \in L_1$ so that r_ℓ, s_ℓ are both elliptic and let $N = \sum_{\ell \in L_1} \mathsf{d}_{\mathcal{T}}(v_\ell, w_\ell)$. Subivide C_0 and its translates so that each hyperplane is replaced by $N+1$ parallel copies, indexed by the multiset $\mathcal{E}$ of edges e so that e separates v_ℓ, w_ℓ for some $\ell \in L_1'$ (it is a multiset since we count an edge e multiple times according to the pairs v_ℓ, w_ℓ that it separates). For each $e \in \mathcal{E}$, let $\mathcal{H}_e$ be a set of walls, one parallel to each wall in $\mathcal{H}$, so that $\mathcal{H}_e \cap \mathcal{H}_{e'} = \emptyset$.

We will modify the walls in $\bigcup_{e \in \mathcal{E}} \mathcal{H}_e$ and then replace $\mathcal{H}$ with $\mathcal{H} \sqcup \left(\bigcup_{e \in \mathcal{E}} \mathcal{H}_e\right)$ to produce a new set of walls with the desired properties.

There are two cases:

- if $\mathsf{d}_{\mathcal{T}}(v_\ell, w_\ell) = 1$, then hypothesis (5) provides $\ell' \in L_0$ so that $r_{\ell'}$ stabilizes the edge e joining v_ℓ, w_ℓ and $s_{\ell'} = r_\ell$. By hypothesis (3), $r_{\ell'}$ and $s_{\ell'}$ are not parallel. Either $\mathcal{H}_e$ contains a wall cutting r_ℓ but not s_ℓ or vice versa, and we do nothing to $\mathcal{H}_e$, or each wall cutting r_ℓ cuts s_ℓ and vice versa. In the latter case, since r_ℓ and $r_{\ell'}$ are not parallel, there exists a hyperplane $\widetilde{H}$ of the translate of C_0 corresponding to v_ℓ such that $\widetilde{H}$ cuts $r_{\ell'}$ but not r_ℓ.

  Let $\widetilde{U}$ be the wall containing $\widetilde{H}$, which intersects the translate of C_0 over w_ℓ in a hyperplane $\widetilde{H}'$. Let $\widetilde{U}'$ be a wall cutting r_ℓ and s_ℓ. Then, equivariantly modifying the arcs of $\widetilde{U}, \widetilde{U}'$ as shown in Figure 5, we replace $\widetilde{U}, \widetilde{U}'$ by a new pair of walls, one cutting r_ℓ and one cutting s_ℓ.

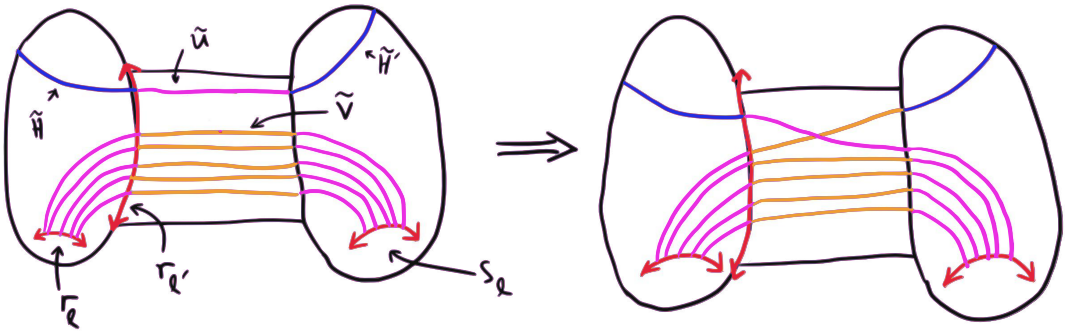

FIGURE 5

- if $\mathsf{d}_{\mathcal{T}}(v_\ell, w_\ell) \geq 2$, then hypothesis (5) provides $\ell' \in L_0$ and adjacent edges e, e' separating v_ℓ, w_ℓ, so that $r_{\ell'}$ stabilizes e and $s_{\ell'}$ stabilizes e'. Hypothesis (3) ensures that $r_{\ell'}, s_{\ell'}$ are not parallel, i.e. there is a hyperplane $\widetilde{H}$ of the translate of C_0 corresponding to $e \cap e'$ so that $\widetilde{H}$ cuts exactly one of $r_{\ell'}, s_{\ell'}$; assume it cuts $r_{\ell'}$. Let $\widetilde{U}$ be the wall of $\mathcal{H}_e$ containing $\widetilde{H}$.
  Either some wall in $\mathcal{H}$ cuts r_ℓ but not s_ℓ, or vice versa, and we do not modify $\mathcal{H}_{e''}$ for any edge e'' associated to ℓ, or every wall cutting r_ℓ cuts s_ℓ and vice versa. Let $\widetilde{U}'$ be such a wall. Then, as shown in Figure 6, we modify the arcs of $\widetilde{U}, \widetilde{U}'$ to produce two new walls, one cutting r_ℓ but not s_ℓ and the other cutting s_ℓ but not r_ℓ. We replace $\widetilde{U}, \widetilde{U}'$ in $\mathcal{H}_e$ by these new walls.

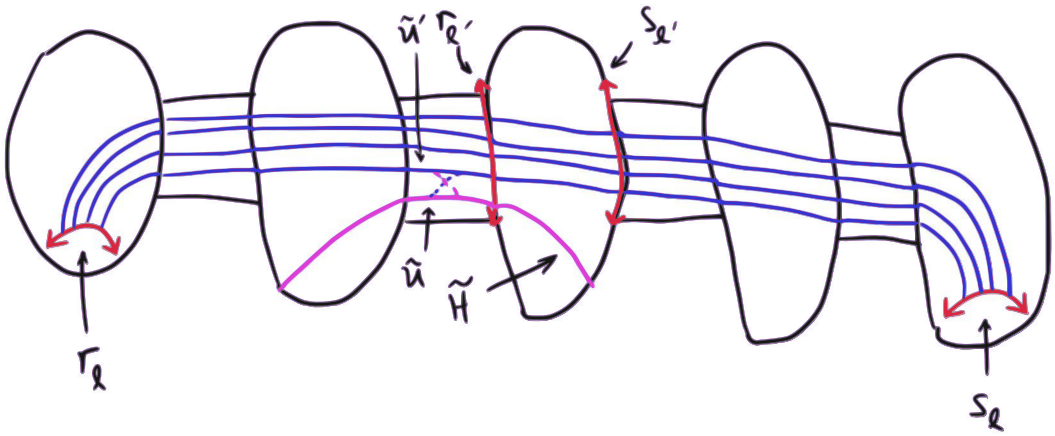

FIGURE 6

There is also a collection $\mathcal{V}$ of *vertical* walls, which are elevations to $\widehat{D}$ of $B \times \{0\}$. Let $\mathcal{W} = \mathcal{V} \cup \mathcal{H}$ and let $\widetilde{\mathcal{W}}$ be the G'–invariant set of walls in $\widetilde{D}$ which are lifts of the elements of $\mathcal{W}$. Given $g \in G'$, let $\|g\|_{\mathcal{W}}$ be the number of $\langle g \rangle$–orbits of walls in $\widetilde{\mathcal{W}}$ cutting g. By Lemma 5.5, every $g \in G' - \{1\}$ satisfies $\|g\|_{\mathcal{W}} > 0$. Let C_1' be the cube complex dual to the G'–wallspace $(\widetilde{D}, \widetilde{\mathcal{W}})$, on which we have just shown that G' acts freely.

Checking $\widetilde{\mathcal{W}}$–non-parallelism in L_0: For $\ell \in L_0$, the elliptic elements r_ℓ, s_ℓ continue to be non-parallel with respect to the set $\widetilde{\mathcal{W}}$ of walls, since they fix a common vertex-space, by Lemma 5.4.

Checking $\widetilde{\mathcal{W}}$–non-parallelism in L_1: Let $\ell \in L_1$. We claim that there is a wall $\widetilde{W} \in \widetilde{\mathcal{W}}$ that cuts r_ℓ but not s_ℓ or vice versa. We divide into four cases according to how r_ℓ, s_ℓ act on $\mathcal{T}$:

(1) r_ℓ, s_ℓ have a common axis $\widetilde{M}$ in $\mathcal{T}$.
(2) r_ℓ, s_ℓ have distinct axes $\widetilde{M}_r, \widetilde{M}_s$.
(3) r_ℓ has an axis $\widetilde{M}$ and s_ℓ is elliptic (or vice versa).
(4) r_ℓ and s_ℓ fix vertices v_r, v_s (which must be distinct by the definition of L_1).

Case (1) is ruled out by Lemma 2.17 by hypothesis (6) and the fact that r_ℓ, s_ℓ do not lie in the same maximal cyclic subgroup.

In case (2), there is a vertical wall that cuts r_ℓ but not s_ℓ.

In case (3), there is a vertical wall that cuts r_ℓ but does not intersect any edge space stabilized by s_ℓ and hence does not cut s_ℓ. Indeed, suppose that every vertex space in $\widetilde{M}$ is stabilized by s_ℓ. Then there exists a conjugate u of t so that $s_\ell \in \cap_{n\in\mathbb{Z}}\langle r_\ell^n u r_\ell^{-n}\rangle$. By Lemma 2.17, r_ℓ commutes with a positive power of u, whence, by hypothesis, r_ℓ is conjugate into $Z_{G'}(t)$, contradicting hypothesis (6). In case (4), the wall $\widetilde{W}_\ell$ constructed above cuts r_ℓ but not s_ℓ or vice versa.

Turning horizontal walls away from $\{q_k t\}$: For each $k \in K$, let $\mathcal{C}_k \subseteq \widetilde{\mathcal{W}}$ be a set of all walls $\widetilde{U}$ with the following properties:

- $\widetilde{U}$ cuts $q_k t$;
- there exists $p \in \{p_j : j \in J\} \sqcup \{1\}$ and $g \in G'$ such that $\widetilde{U}$ cuts $g(pt)g^{-1}$.

Note that each $\widetilde{U} \in \mathcal{C}_k$ chops $q_k t$. Let $\widetilde{\mathcal{W}}'$ be the collection of walls obtained from $\widetilde{\mathcal{W}}$ by adding an additional parallel copy of each wall in $G'(\cup_{k\in K}\mathcal{C}_k)$. Let $\widetilde{X}$ be the cube complex dual to $\widetilde{\mathcal{W}}'$.

Given a wall $\widetilde{W} \in \widetilde{\mathcal{W}}'$, we use the same notation $\widetilde{W}$ for the corresponding hyperplane of $\widetilde{X}$, and for each $k \in K$, we denote by $\mathcal{C}_k'$ the set of hyperplanes corresponding to walls in $\mathcal{C}_k$, so that the elements of $\mathcal{C}_k'$ come in parallel pairs.

Since each edge-space in $\widetilde{D}$ was included as a (vertical) wall, $\widetilde{X}$ decomposes as a tree of spaces where the underlying tree is $\mathcal{T}$ and the edge spaces are hyperplanes. For each $k \in K$, let $\widetilde{Y}_k$ be the preimage of the $q_k t$–axis in $\mathcal{T}$ under $\widetilde{X} \to \mathcal{T}$. Note that $\widetilde{Y}_k$ is a convex subcomplex containing the convex hull of every combinatorial $q_k t$–axis in $\widetilde{X}$.

Let $X = G \setminus \widetilde{X}$. By convexity, the inclusion $\widetilde{Y}_k \hookrightarrow \widetilde{X}$ descends to a local isometry $\widetilde{Y}_k \to X$, and, letting $\mathcal{Y} = \{\widetilde{Y}_k : k \in K\}$, we consider the cubical presentation $X^* = \langle X \mid \mathcal{Y}\rangle$. Lemma 5.6 shows that X^* satisfies the $C'(\frac{1}{12})$ condition. By Lemma 5.8, each $\widetilde{Y}_k$ is endowed with a wallspace structure making X^* a $B(6)$ cubical presentation.

We now verify that X^* satisfies the remaining hypothesis of Theorem 2.15. Specifikally, fix k and suppose that κ is a geodesik of $\widetilde{Y}_k$ with initial and terminal 1–kubes respektively dual to hyperplanes H_p, H_q. Consider distinct hyperplanes W, W' lying in the same wall of $\widetilde{Y}$. Suppose κ traverses W, and W' is proximate to H_q, i.e. there is a path P and a 1–cube e_q dual to H_q so that $e_q P$ joins H_q to W and either $e_q P$ is a piece-path or P is trivial. Without loss of generality, we may assume that the projection of κ to $\mathcal{T}$ is such that the hyperplanes are all ordered as in Figure 7.

Note that $W' = (q_k t)^M W$, by construction of the wallspace described in the proof of Lemma 5.8, where M has the property that no path starting on W and ending on W' is a concatenation of fewer than 7 piece-paths. We now add an additional constraint on M: choose $n_0 \geq 0$ so that $(q_k t)^{n_0}$ has translation length greater than M_0 in $\mathcal{T}$, where M_0 is an upper bound on the diameter of the projection of any piece-path to $\mathcal{T}$, exhibited in the proof of Lemma 5.8. Then require $M > 100n_0$.

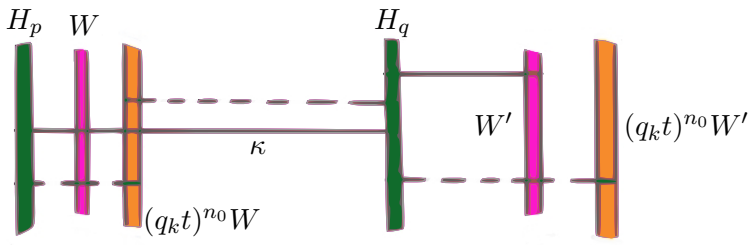

FIGURE 7. The middle dashed path is not a concatenation of fewer than 98 piece-paths, and the left and right dashed paths are not piece-paths. Thus the orange wall separates H_p, H_q and is not proximate to either.

Consider the $\widetilde{Y}_k$–wall consisting of the hyperplanes $(q_kt)^{n_0}W, (q_kt)^{n_0}W'$, shown in Figure 7. Then $(q_kt)^{n_0}W$ intersects κ between W and H_q and cannot be proximate to either W or H_q, and in particular $(q_kt)^{n_0}W$ separates H_p, H_q. On the other hand, $(q_kt)^{n_0}W'$ cannot be proximate to W' and is hence separated from H_q by W'. Thus this wall separates H_p, H_q and is not proximate to H_q.

Hence Theorem 2.15 provides a free action of G' on a CAT(0) cube complex $\widetilde{X}_1$ with the following properties:

- $\widetilde{X}_1$ splits as a tree of spaces, with the same underlying tree $\mathcal{T}$ as before, with edge-spaces that are hyperplanes corresponding to original vertical walls.
- Every hyperplane $\widetilde{H}$ of $\widetilde{X}$ not belonging to $G'\left(\cup_{k\in K}\mathcal{C}'_k\right)$ corresponds to a unique hyperplane $\widetilde{H}_1$ of $\widetilde{X}_1$, and $\widetilde{H}_1$ cuts $g \in G'$ in $\widetilde{X}_1$ if and only if $\widetilde{H}$ cuts g in $\widetilde{X}$. Indeed, each such $\widetilde{H}$ is an entire wall in $\widetilde{X}^*$.
- $\|p_jt\|_{\widetilde{X}_1} = \|p_jt\|_{\widetilde{X}} = \|p_jt\|_{C_0} = \|t\|_{C_0} = \|t\|_{\widetilde{X}} = \|t\|_{\widetilde{X}_1}$ for all $j \in J$. Indeed, in passing from C_0 to $\widetilde{X}$ (i.e. forming the graph of spaces and constructing $\widetilde{\mathcal{W}}$), we did not affect the translation lengths of t, p_jt by Lemma 5.4, while in passing from $\widetilde{X}$ to $\widetilde{X}_1$ (i.e. in turning the walls in the various $\mathcal{C}_k$), we did not affect the translation length of any elliptic element. Indeed, walls in $\widetilde{X}^*$ have connected intersection with vertex spaces by Theorem 2.16 and the fact that hyperplanes of any $\widetilde{Y}_k$ that lie in the same wall of $\widetilde{Y}_k$ do not intersect the same vertex space.
- Each q_kt is cut in $\widetilde{X}_1$ by a hyperplane corresponding to a vertical wall.
- For each $p \in \{p_j : j \in J\} \cup \{1\}$ and each $g \in G'$ and $k \in K$, no hyperplane of $\widetilde{X}_1$ cuts both pt and $g(q_kt)g^{-1}$.
- If $\alpha \in G'$ is elliptic, and $\widetilde{H}$ is a hyperplane of $\widetilde{X}$, then the hyperplane $\widetilde{W}$ of $\widetilde{X}_1$ corresponding to the wall in $\widetilde{X}^*$ containing $\widetilde{H}$ cuts α if and only if $\widetilde{H}$ cuts α.

The fourth assertion holds because the equivalence class of hyperplanes containing a vertical wall was a singleton, since these walls do not cut any conjugate of any q_kt. The fifth assertion holds because, if such a hyperplane existed, then some p_jt–cutting [resp. t–cutting], $g(q_kt)g^{-1}$–chopping hyperplane was paired with some other $g(q_kt)g^{-1}$–chopping hyperplane, so the resulting wall does not cut $g(q_kt)g^{-1}$. Indeed, by Theorem 2.16, this wall intersects $\widetilde{Y}_k$ in exactly this pair of hyperplanes.

Turning vertical walls: We now pass from $\widetilde{X}_1$ to a new cube complex C'_1 by constructing a G'–invariant wallspace on $\widetilde{X}$ with the following properties:

(1) Every hyperplane of $\widetilde{X}$ that is non-vertical continues to be a wall.
(2) For each $k \in K$, there exists a wall $\widetilde{V}_k$ that cuts $q_k t$ but does not cut any conjugate of $q_{k'}t$ for $k' \neq k$ and does not cut any elliptic element. In particular, $\widetilde{V}_k$ does not cut any conjugate of t or $p_j t$ for $j \in J$.

$\widetilde{X}_1$ decomposes as a tree of cube complexes with underlying tree $\mathcal{T}$ and hyperplane edge spaces. Subdivide so that $\widetilde{X}_1$ includes $2|K|+1$ parallel copies of each vertical hyperplane (i.e. edge space). For each $k \in K$, let $\widetilde{Z}_k$ be the preimage of the $q_k t$–axis in $\mathcal{T}$ under the map $\widetilde{X}_1 \to \mathcal{T}$, and observe that $\widetilde{Z}_k$ is a convex subcomplex of $\widetilde{X}_1$. Let $X_1 = G' \setminus \widetilde{X}_1$, and consider the cubical presentation $X_1^* = \langle X_1 \mid \{\widetilde{Z}_k\}_{k\in K}\rangle$. Lemma 5.6 shows that X_1^* is a $C'(\frac{1}{12})$ presentation.

We now declare the wallspace structure on $\widetilde{Z}_k$ as follows. For each edge e of $\mathcal{T}$, equivariantly label its $2|K|+1$ vertical hyperplanes $\{\widetilde{V}_i^{\pm} : i \in K\} \cup \{\widetilde{V}\}$. Let $k \in K$. We define an equivalence relation $\sim_k$ on the hyperplanes of $\widetilde{Z}_k$ as follows. For each e and each edge-space of $\widetilde{Z}_k$ mapping to a translate ge of e, and some $N > 0$ to be determined, declare $(q_k t)^N g\widetilde{V}_k^+ \sim_k g\widetilde{V}_k^-$. Every other equivalence class is a singleton. In particular, the $\sim_k$–classes of $\widetilde{V}, \widetilde{V}_i^+, \widetilde{V}_i^-$ are singletons when $i \neq k$.

Each wall is the union of the one or two hyperplanes in a $\sim_k$–class. The variation of Lemma 5.8 obtained by replacing $\widetilde{Y}_i$ with $\widetilde{Z}_k$ shows that, if N is chosen sufficiently large, X_1^* is $B(6)$.

Let H_p, H_q be hyperplanes intersecting the initial and terminal edges of a geodesic κ in $\widetilde{Z}_k$, and let $\widetilde{W}, \widetilde{W}'$ be $\sim_k$–equivalent hyperplanes such that W intersects κ and W' is proximate to H_q. Then, provided N was chosen sufficiently large, there is some edge-space separating H_p, H_q and containing a translate of $\widetilde{V}$ that separates H_p, H_q and is not proximate to either. Theorem 2.15 now produces a free action of G' on a CAT(0) cube complex C_1' dual to the wallspace induced on $\widetilde{X}_1^*$ by the new walls. While C_1' no longer decomposes as a tree of spaces with underlying tree $\mathcal{T}$, the new dual cube complex C_1' has the properties (1),(2) above.

Subdividing to equalize: We are now in the following situation: G' acts freely on the CAT(0) cube complex C_1' and we have the following:

- $\|t\|_{C_1'} = \|p_j t\|_{C_1'}$ for all $j \in J$.
- For each $p \in \{p_j : j \in J\} \cup \{1\}$ and each $g \in G'$ and $k \in K$, no hyperplane of C_1' cuts both pt and $g(q_k t)g^{-1}$.
- For each $k \in K$, there exists a hyperplane $\widetilde{V}_k$ that cuts $q_k t$ but does not cut any conjugate of $q_{k'}t$ for $k' \neq k$ and does not cut any elliptic element (in particular, $\widetilde{V}_k$ does not cut any conjugate of t or $p_j t, j \in J$).

Given a hyperplane $\widetilde{H}$, the *n–fold subdivision along $\widetilde{H}$* is the cubical subdivision in which $\widetilde{H}$ and each of its G'–translates is replaced by $n+1$ parallel copies of itself. With respect to the n–fold subdivision along $\widetilde{H}$, the translation length of each $g \in G'$ is increased by na, where a is the number of $\langle g\rangle$–orbits of G'–translates of $\widetilde{H}$ that cut g. Thus it does not change if $n = 0$ or no G'–translate of $\widetilde{H}$ cuts g. We will now subdivide C_1' along various hyperplanes to obtain the final cube complex C_1, in which each of $p_j t, q_k t, t, j \in J, k \in K$ have the same translation length.

For each $k \in K$, observe that subdivision along $\widetilde{V}_k$ does not affect the length of any conjugate of any $q_{k'}t, k' \neq k$ or $p_j t$ or t. Choose M so that

$$M \geq \max\left\{ \|q_k t\|_{C_1'} : k \in K \right\}$$

and $M = m\|t\|_{C_1'}$ for some $m \in \mathbb{N}$. Let $n_k = M - \|q_k t\|_{C_1'}$ and let $m = M - \|t\|_{C_1'}$. We now subdivide C_1' as follows. First n_k–fold subdivide along each $\widetilde{V}_k$. Second, $(m-1)$–fold subdivide along each hyperplane cutting t or $p_j t$. Note that this subdivision does not affect the lengths of the conjugates of the various $q_k t$. We conclude that $\|q_k t\|_{C_1} = M$ for each $k \in K$, and $\|t\|_{C_1} = \|p_j t\|_{C_1} = m\|t\|_{C_1'} = M$.

Persistence of non-parallelism in L_0, L_1 after turns and subdivisions: To complete the proof, we verify the following claim.

Claim 2. *Let $\alpha, \beta \in G'$. Suppose that each of α, β is either elliptic or not conjugate into $Z_{G'}(t)$. Then if a hyperplane of $\widetilde{X}$ cuts α but not β, then some hyperplane of C_1 cuts exactly one of α, β.*

Proof of Claim 2. First we show that non-parallelism of α, β persists when passing from $\widetilde{X}$ to $\widetilde{X}_1$. Each of the new walls used to build $\widetilde{X}_1$ intersects each vertex space of $\widetilde{X}$ in at most one hyperplane. Suppose α is elliptic and cut by a hyperplane $\widetilde{H}$ of $\widetilde{X}$ not cutting β. If β is hyperbolic, then it is cut by a vertical wall, and we are done. Suppose that β is elliptic. By the construction of the new walls in $\widetilde{X}$, a hyperplane of $\widetilde{X}$ cuts an elliptic element if and only if the new wall containing it cuts that element, so the hyperplane of $\widetilde{X}_1$ corresponding to the new wall containing $\widetilde{H}$ cuts α but not β.

If α is hyperbolic, then there is a vertical wall cutting exactly one of α but not β, or both are hyperbolic and have the same axis in $\mathcal{T}$. In this case, Lemma 2.17 and the fact that α is not conjugate into $Z_{G'}(t)$ imply that α, β belong to the same maximal cyclic subgroup, which means that it was impossible for some $\widetilde{H}$ to cut α but not β in $\widetilde{X}$, a contradiction.

Second, we show that non-parallelism persists in passing from $\widetilde{X}_1$ to C_1' (and hence to C_1, since C_1 is obtained from C_1' by iterated subdivision). Suppose first that α is hyperbolic in $\widetilde{X}_1$ and a vertical hyperplane $\widetilde{H}$ cuts α but not β (or vice versa). Then, after the initial subdivision, there are $2|K|+1$ vertical hyperplanes of $\widetilde{X}_1$ cutting α but not β, one of which continues to be a wall in the new wallspace used to form C_1'; the corresponding hyperplane of C_1' therefore cuts α but not β.

Now suppose that α, β are hyperbolic in $\widetilde{X}_1$ and a non-vertical hyperplane $\widetilde{H}$ cuts α but not β. Then the equivalence class containing $\widetilde{H}$ is a singleton, i.e. the new wall containing $\widetilde{H}$ cuts α but not β.

The claim holds in the case where α, β are both elliptic since a hyperplane of $\widetilde{X}_1$ cuts an elliptic element if and only if the new wall in $\widetilde{X}_1$ containing it does. Indeed, they are the same. ■

Claim 2 implies the statement about elements r_ℓ, s_ℓ with $\ell \in L_0 \cup L_1$ in view of the above discussion of $\widetilde{X}$–non-parallelism, and the proof is complete. □

Lemma 5.4. *Let $\widetilde{W}$ be a horizontal wall of $\widetilde{D}$, and let $\widetilde{W}'$ be a vertical wall. Then $|\widetilde{W} \cap \widetilde{W}'| \leq 1$, and hence the intersection of $\widetilde{W}$ with a translate of $\widetilde{C}_0$ is either empty or consists of a single hyperplane.*

Proof. This holds since $\widetilde{W}$ is a tree of spaces whose underlying tree embeds in $\mathcal{T}$. □

Lemma 5.5. *Each hyperbolic $g \in G'$ is cut by a vertical wall $\widetilde{W}'$ of $\widetilde{D}$. Each nontrivial (elliptic) $g \in \pi_1 D$ is cut by a horizontal wall $\widetilde{W}$.*

Proof. The first claim is immediate. The second claim holds by Lemma 5.4. □

Lemma 5.6. *Let $X^* = \langle X \mid \{\widetilde{Y}_k : k \in K\}\rangle$ be a cubical presentation with each $\widetilde{Y}_k$ simply-connected. Then X^* satisfies the $C'(\frac{1}{12})$ condition.*

Proof. Each $\widetilde{Y}_k$ is simply–connected, so there are no essential paths in $\widetilde{Y}_k$. □

Remark 5.7. In fact, the proof of Lemma 5.6 shows that $C'(\alpha)$ holds for all $\alpha > 0$.

Lemma 5.8. *Each $\widetilde{Y}_i$ admits a wallspace structure so that $\langle X \mid \mathcal{Y}\rangle$ satisfies the $B(6)$ condition.*

Proof. Recall that for each k, we have a set $\mathcal{C}'_k$ of hyperplanes in $\widetilde{X}$ so that $\mathcal{C}'_k$ consists of at most one hyperplane from each $\langle q_k t\rangle$–orbit of hyperplanes cutting $q_k t$, but does not include any hyperplanes corresponding to vertical walls in $\widetilde{\mathcal{W}}$. Recall that the elements of $\mathcal{C}'_k$ come in osculating parallel pairs, so that we can write $\mathcal{C}'_k = \{W_1^+, W_1^-, \dots, W_n^+, W_n^-\}$ where $W_i^\pm$ are parallel and osculate.

Fix $N > 0$ and choose $M > 0$ so that for each i, we have that $(q_k t)^M W_i^\pm$ and $W_i^\pm$ are disjoint and separated by at least N edge spaces. We now declare $(q_k t)^n W_i^\pm \sim_k (q_k t)^{n+M} W_i^\mp$ for each i and each $n \in \mathbb{Z}$. For any other hyperplanes U, V of $\widetilde{Y}_k$, we declare $U \sim_k V$ if and only if $U = V$. The equivalence relation $\sim_k$ has the property that $\sim_k$–equivalent hyperplanes do not cross or osculate, and each $\sim_k$–class partitions $\widetilde{Y}_k$ into hyperplanes in the obvious way. Clearly $\sim_k$ is $\langle q_k t\rangle$–invariant, and hence $\mathrm{Aut}(\widetilde{Y}_k \to X)$–invariant by Lemma 2.17. It thus remains to verify that, when N is sufficiently large, condition (4) of Definition 2.13 is satisfied.

By Lemma 2.17, there exists $M_0 \geq 1$ such that any piece-path in a cone-piece projects to a path in $\mathcal{T}$ whose endpoints lie at distance at most M_0 in $\mathcal{T}$. Any hyperplane $\widetilde{H}$ of $\widetilde{X}$ that does not intersect $\widetilde{Y}_k$ but whose carrier intersects $\widetilde{Y}_k$ has the property that $N(\widetilde{H}) \cap \widetilde{Y}_k$ is contained in a single vertex-space of $\widetilde{X}$.

Let $P \to \widetilde{Y}_k$ be a path starting and ending on $N(W)$ for some wall W. If W is formed from a $\sim_k$–class with a single element, i.e. W is a hyperplane, then P is homotopic into $N(W)$ since $\widetilde{Y}_i$ is CAT(0). Otherwise, W is the union of two hyperplanes with the property that any path in $\widetilde{Y}_i$ joining them must traverse at least N edge-spaces, and hence either P starts and ends on one of these hyperplanes (and is thus homotopic into $N(W)$) or P cannot be the concatenation of fewer than $\frac{N}{M_0}$ piece-paths. By taking $N > 7M_0$, we ensure that the desired condition is satisfied. (In fact, by choosing N sufficiently large, we can ensure that the following holds for any specified K: if $P \to Y_i$ is a path that is the concatenation of at most K piece-paths and P starts and ends on the carrier $N(U)$ of a wall then P is path-homotopic into $N(U)$.) □

5.1. **A finite cover where $q_k t$–cycles embed.**

Lemma 5.9. *Let $G = A*_{\langle b\rangle^e = \langle c\rangle}$ be an HNN extension with $\langle b\rangle \cong \mathbb{Z}$. Suppose $\langle b\rangle$ is separable in G. Let $g_1, \dots, g_k \in G$ act hyperbolically on the Bass-Serre tree so that each g_i has minimal translation length in its conjugacy class. There exists a finite-index normal subgroup $G' \leq G$ whose induced splitting as a graph Γ' of groups has the following property: for each i, let $\langle g_i'\rangle = \langle g_i\rangle \cap G'$. The normal form for g_i' projects to an embedded cycle in Γ' for each i, in the sense that all stable letters in g_i' map to distinct edges in Γ'. Moreover, we can assume Γ' is simplicial.*

Proof. Since we can take an intersection of conjugates of finite-index subgroups constructed for each i, it suffices to find a finite-index (not necessarily normal) $G_i \leq G$ with the desired property for each particular g_i. The element g_i has a normal form $a_0 e^\pm \cdots a_{n-1} e^\pm a_n$ whose characteristic property is that it does not (even cyclically) contain a subword of the form $eb^p e^{-1}$ or $e^{-1}c^p e$, and note that there is at least one $e^{\pm 1}$ in this expression since g_i is hyperbolic. For each j such that $ea_j e^{-1}$ appears in the normal form, we let B_j be a finite-index subgroup of G with $\langle b\rangle \leq B_j$ but $a_j \notin B_j$. Likewise, when $e^{-1}a_j e$ appears, let C_j be a finite-index subgroup of G with $\langle c\rangle \leq C_j$ but $a_j \notin C_j$. Let G_i be the intersection of the various B_j, C_j arising above. Note that the normal

form for any power g_i^q projects to an immersed path in the underlying graph Γ_i of the induced splitting, in the sense that successive stable letters project to distinct edges. Separability of cyclic subgroups of $\pi_1\Gamma_i$ allows us to produce a finite cover $\Gamma_i' \to \Gamma_i$ where this path embeds; we let G_i' be the resulting finite-index subgroup of G_i. The "moreover" clause follows by passing to the cover associated to $H_1(\Gamma_i', \mathbb{Z}_3)$. □

Lemma 5.9, immediately yields:

Corollary 5.10. *Let $G' = G*_{\langle t\rangle^e = \langle p_1 t\rangle}$ and $\{q_k t\}$ be as in Theorem 5.2. Then there exists a finite-index normal subgroup $G'' \leq G'$ so that, in the induced splitting of G'' as a simplicial graph Γ'' of groups, the normal form for a generator of $\langle q_k t\rangle \cap G''$ projects to an embedded cycle in Γ for each k.*

5.2. **Using exits to build turnable walls in the finite cover.** Let $\Gamma = (\mathcal{V}, \mathcal{E})$ be a finite connected simplicial graph, and let Y be a space splitting as a graph of spaces with the following structure. For each $v \in \mathcal{V}$, the associated vertex space Y_v is a nonpositively-curved cube complex with a finite set $\mathcal{H}_v$ of immersed hyperplanes. For each edge e, the associated edge space C_e is a cylinder with attaching maps $f_{e,v}, f_{e,w} : \mathbb{S}^1 \to Y_v, Y_w$, where v, w are the vertices of e.

We also require that each $f_{e,v}$ has the property that for each $H \in \mathcal{H}_v$, the set $f_{e,v}^{-1}(H)$ is finite and, moreover, $f_{e,v}$ has been chosen in its homotopy class to minimize $|f_{e,v}^{-1}(H)|$ for each H. Moreover, for each hyperplane H of Y_v, the intersections of H with the various incident edge-spaces are disjoint.

We impose some additional constraints, which are satisfied for the graph of spaces in Theorem 5.2. First, for each $e \in \mathcal{E}$ with endpoints v, w,

$$\textbf{(I)} \qquad \sum_{H\in\mathcal{H}_v} |f_{e,v}^{-1}(H)| = \sum_{H\in\mathcal{H}_w} |f_{e,w}^{-1}(H)|,$$

i.e. each edge space has the same total number of intersections with hyperplanes in its initial and terminal vertex spaces. Second, Γ contains a particular finite set $\mathcal{C}$ of embedded cycles, such that for each $\sigma \in \mathcal{C}$, there exist $v_\sigma \in \mathcal{V}, H_\sigma \in \mathcal{H}_{v_\sigma}$ such that

$$\textbf{(II)} \qquad |f_{e,v_\sigma}^{-1}(H_\sigma)| = 0$$

for **exactly one edge** e of σ incident to v_σ. The hyperplane H_σ is an *exit* for σ and Y_{v_σ} is an *exit vertex space* for σ.

For each hyperplane H of Y_v and $n \in \mathbb{N}$, let nH be the disjoint union of the n parallel copies of H which would be hyperplanes in the $(n-1)$–fold subdivision of Y_v. Let $W_n = \bigsqcup_{v\in\mathcal{V}} \bigsqcup_{H\in\mathcal{H}_v} nH$. Then we have an immersion $I_n : W_n \to Y$ defined as follows: replace each Y_v with the $(n-1)^{th}$ cubical subdivision, and map the n copies of each $H \in \bigcup_{v\in\mathcal{V}} \mathcal{H}_v$ bijectively to the parallel copies of $H \to Y_v$. For each $e \in \mathcal{E}$ and vertex v of e, and $H \in \mathcal{H}_v$, consider the set $f_{e,v}^{-1}(nH)$ of intersection points of $f_{e,v}$ with the n parallel copies of H. Observe that

$$\sum_{H\in\mathcal{H}_v} |f_{e,v}^{-1}(nH)| = \sum_{H'\in\mathcal{H}_{e^-}} |f_{e,w}^{-1}(nH')|,$$

where w is the other edge of e, by condition **(I)**. A W_n–*matching* is a set of bijections

$$\left\{ b_e \colon \bigcup_{H\in\mathcal{H}_v} f_{e,v}^{-1}(nH) \to \bigcup_{H'\in\mathcal{H}_w} f_{e,w}^{-1}(nH) \right\}_{e\in\mathcal{E}}$$

Given a W_n–matching, for each $e \in \mathcal{E}$ and each $a \in \bigcup_{H \in \mathcal{H}_v} f_{e,v}^{-1}(nH)$, choose an properly embedded arc α in $C_e \times [0,1] \to Y$ so that the image of α in Y joins a to $b_e(a)$. Then we have an immersed *W_n–matching space* $V \to Y$, depending on $\{b_e : e \in \mathcal{E}\}$ and the chosen arcs, which is a graph of spaces with vertex spaces the components of W_n and edge-spaces the arcs joining each a to $b_e(a)$. The arc joining a to $b_e(a)$ in V is sent by $V \to Y$ homeomorphically to the arc α chosen above.

Lemma 5.11. *Let $V \to Y$ be a W_n–matching space. Then each component C of V is an immersed wall in Y. Moreover, the universal cover $\widetilde{C}$ lifts to a wall in the universal cover $\widetilde{Y}$ intersecting each vertex space in a hyperplane or $\emptyset$ and intersecting each edge space in an embedded arc or $\emptyset$.*

Let $V \to Y$ be a W_n–matching space, with underlying graph Λ_V. The map $V \to Y$ induces a homomorphism $\Lambda_V \to \Gamma$ of graphs, sending each vertex, corresponding to an immersed hyperplane in some Y_v, to the vertex v, and each edge, corresponding to an arc in some $C_e \times [0,1]$, to the edge e. The space V *avoids* $\mathcal{C}$ if no cycle $\hat{\sigma}$ of Λ_V has the property that $\Lambda_V \to \Gamma$ restricts to a covering map $\hat{\sigma} \to \sigma$ for some $\sigma \in \mathcal{C}$. In the latter case, we call $\hat{\sigma}$ a *bad σ–cycle* in Λ_V.

Theorem 5.12 (Building $\mathcal{C}$–avoiding matching spaces)**.** *There exists $n \in \mathbb{N}$ and a W_n–matching space $V \to Y$ that avoids $\mathcal{C}$. Moreover, for each $v \in \mathcal{V}$ and $H \in \mathcal{H}_v$, either $nH \subset V$ or $f_{e,v}^{-1}(nH) = \emptyset$ for each e incident to v.*

Proof. The "moreover" clause follows from the definition of a W_n–matching space. Hence it suffices to prove the first assertion.

For $r \in \mathbb{N}$, choose a homomorphism $\phi : \mathsf{H}_1(\Gamma) \to \mathbb{Z}$ such that $\phi([\sigma]) \neq 0$ for each $\sigma \in \mathcal{C}$. For $r \geq 2$, let ϕ_r be the composition $\mathsf{H}_1(\Gamma) \xrightarrow{\phi} \mathbb{Z} \to \mathbb{Z}_r$ and let Γ_r be the cover associated to $\ker\left(\pi_1\Gamma \to \mathsf{H}_1(\Gamma) \xrightarrow{\phi_r} \mathbb{Z}_r\right)$. Let $\widehat{\mathcal{C}}$ be the set of components of preimages of cycles in $\mathcal{C}$. Let $\widehat{Y}_r \to Y$ be the corresponding cover of Y.

We will consider $\widehat{Y}_r$ for all sufficiently large primes r. This ensures that for each $\sigma \in \mathcal{C}$, the preimage of σ in $\widehat{Y}_r$ is a cycle $\hat{\sigma}$ of length $r|\sigma|$ covering σ, so the map $\hat{\sigma} \mapsto \sigma$ is a bijection $\widehat{\mathcal{C}} \to \mathcal{C}$.

Observe that any W_1–matching of the hyperplanes in $\widehat{Y}_r$ determines a W_r–matching of the hyperplanes of Y. Indeed, if $V \to \widehat{Y}_r$ is a W_1–matching space, then the composition $V \to \widehat{Y}_r \to Y$ can be perturbed so that for each hyperplane $H \to Y$ in the image, the r copies of H in V map to disjoint parallel copies of H in Y.

Let $\widehat{\mathcal{E}}$ be the edge set of $\widehat{\Gamma}$ and choose a matching $\{b_e : e \in \widehat{\mathcal{E}}\}$ uniformly randomly among all possible W_1–matchings on $\widehat{Y}_r$. We claim that the probability that some W_1–matching space $V \to \widehat{Y}_r$ associated to $\{b_e\}$ avoids $\widehat{\mathcal{C}}$ tends to 1 as $r \to \infty$.

First, observe that since V is a W_1–matching space, and since each $\hat{\sigma} \in \widehat{\mathcal{C}}$ has length $r|\sigma|$ (where $\hat{\sigma} \to \sigma \in \mathcal{C}$), the matching space V avoids $\widehat{\mathcal{C}}$ if and only if the underlying graph Λ_V of V contains no element of $\widehat{\mathcal{C}}$.

For each $\hat{\sigma} \in \widehat{\mathcal{C}}$, projecting to $\sigma \in \mathcal{C}$, consider the probability $p_{\hat{\sigma}}$ that there is a cycle of length $r|\sigma|$ in Λ_V projecting to $\hat{\sigma}$. Let $v \in \sigma$ be a vertex with the property that Y_v contains an exit hyperplane H_σ for σ, so there is exactly one edge e of σ incident to v with the property that $(f_{e,v})^{-1}(H_\sigma) \neq \emptyset$. Let w be the other endpoint of e. Let $\mathfrak{E}_v \subset \mathcal{H}_v$ be the set of σ–exit hyperplanes in Y_v.

Let

$$\alpha_\sigma(e,v) = \frac{\sum_{H \in \mathfrak{E}_v} |f_{e,v}^{-1}(H)|}{\sum_{H \in \mathcal{H}_v} |f_{e,v}^{-1}(H)|} \in (0,1]$$

be the proportion of hyperplane points of $f_{e,v}(C_e)$ that are exit hyperplane points.

We claim that the probability that Λ_V contains a bad $\hat{\sigma}$–cycle is bounded by $(1-\alpha_\sigma(e,v))^r$. Let $\hat{v}_1,\dots,\hat{v}_r$ be the lifts of v in $\hat{\sigma}$, in cyclic order, and let $\hat{w}_1,\dots,\hat{w}_r$ be the lifts of w, arranged so that $\hat{w}_i,\hat{v}_i$ are joined by a lift $\hat{e}_i$ of e. Suppose that there is a bad $\hat{\sigma}$–cycle in Λ_V. Then for each i, the matching space V contains a hyperplane H_i of $Y_{\hat{w}_i}$ that intersects both of the edge spaces corresponding to the edges of the cycle incident to $\hat{w}_i$, i.e. H_i is a non-exit hyperplane. Moreover, there is an intersection point of H_i with $f_{\hat{e}_i,\hat{w}_i}(C_{\hat{e}_i})$ which is joined by an arc in V to an intersection point of $f_{\hat{e}_i,\hat{v}_i}(C_{\hat{e}_i})$ and a non-exit hyperplane of $Y_{\hat{v}_i}$. The probability of this occurring at the i^{th} stage is $(1-\alpha_\sigma(e,v))$, and the choices are independent at different stages. Hence the probability that the bad $\hat{\sigma}$–cycle exists is at most $(1-\alpha_\sigma(e,v))^r$.

Figure 8 shows a heuristic picture of the generic situation.

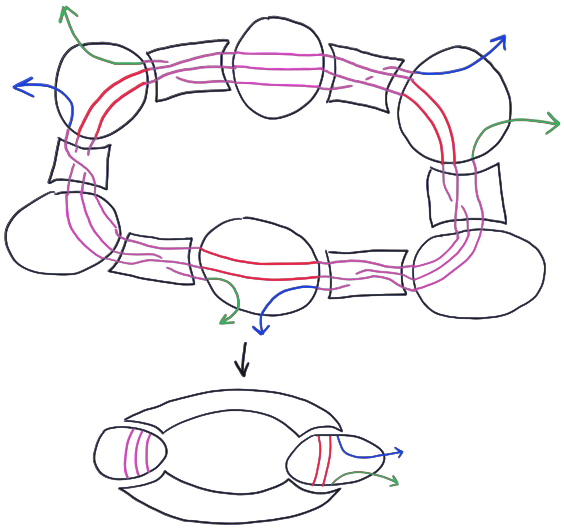

FIGURE 8. The graphs of spaces associated to $\hat{\sigma}$ and σ, and a matching space arising from a typical matching.

Fixing r, for each $\sigma\in\mathcal{C}$ let E_σ be the event that Λ_V contains a bad σ–cycle. (Here we are regarding V as a W_r–matching space mapping to Y, as explained above.) Equivalently, E_σ is the event that Λ_V, regarded as the underlying graph of a W_1–matching space in $\widehat{Y}_r$, contains a bad $\hat{\sigma}$–cycle, since $\hat{\sigma}$ is the preimage of σ under $\widehat{Y}_r\to Y$.

We have just shown that $\mathbb{P}(E_\sigma)\leq(1-\alpha)^r$, where $\alpha=\min_{\sigma\in\mathcal{C}}\{\max_{e,v}\{\alpha_\sigma(e,v)\}\}$, where, for each $\sigma\in\mathcal{C}$, the vertex v varies among those v so that Y_v contains a σ–exit hyperplane and e varies over the edges of σ incident to v. Note that $\alpha>0$. Hence $\mathbb{P}\left(\bigcup_{\sigma\in\mathcal{C}}E_\sigma\right)\leq|\mathcal{C}|(1-\alpha)^r$, which tends to 0 as $r\to\infty$, as required. □

Corollary 5.13 (Turnable walls)**.** *There exists a finite collection $\mathcal{H}$ of immersed walls $W\to Y$ with the following properties:*

1. *each W is a component of a W_N–matching space for some N;*
2. *for each $\sigma\in\mathcal{C}$ and each $g\in\pi_1Y$ projecting to σ, and each $W\in\mathcal{H}$, if a lift $\widetilde{W}$ of W cuts g, then $\widetilde{W}$ chops g.*

Proof. If $\widetilde{W}$ fails to chop g, then $\{g^n\widetilde{W}\}_{n\in\mathbb{Z}}$ is a set of pairwise-intersecting walls. It follows that W does not avoid $\mathcal{C}$. But by Theorem 5.12, we could have chosen $\mathcal{H}$ so that each wall avoids $\mathcal{C}$. □

5.3. **Connecting Theorem 5.2 to the free-by-$\mathbb{Z}$ case.** We now establish facts needed to apply Theorem 5.2 in the proof of Theorem A.

Lemma 5.14. *Let $G = F \rtimes \langle t \rangle$, let $f \in F$, let $H = \langle ft \rangle$. Then $H \leq G$ is separable.*

Proof. The subgroup $\langle ft \rangle$ is a retract of G for each $f \in F$, and a retract of a residually finite group is separable. □

Lemma 5.15. *Let F be a finite-rank free group, let $\Phi : F \to F$ be a polynomially growing automorphism with an improved relative train track (IRTTM) representative and let $G = \langle F, t \mid \{tft^{-1} = \Phi(f) : f \in F\} \rangle$. Suppose that $gt^p g^{-1} = t^q$ for some $g \in G$ and $p, q \neq 0$. Then $gt = tg$.*

Proof. By examining the map to $\langle t \rangle$, we see that $p = q$. To conclude the proof, suppose that $gt^p g^{-1} = t^p$ for some $g \in G - \{1\}$ and $p, q \neq 0$. We can write $g = ft^n$ for some $f \in F$ and $n \in \mathbb{Z}$, so $ft^p f^{-1} = t^p$. Thus $\Phi^p(f) = f$. By the definition of an improved relative train track map [7, Theorem 5.1.5], f is represented by a Nielsen path, i.e. $\Phi(f) = f$. Hence f, t commute. □

References

[1] Algom-Kfir, Y., Hironaka, E., and Rafi, K. Digraphs and cycle polynomials for free-by-cyclic groups. *Geometry & Topology 19*, 2 (2015), 1111–1154.

[2] Algom-Kfir, Y., and Rafi, K. Mapping tori of small dilatation expanding train-track maps. *Topology and its Applications 180* (2015), 44–63.

[3] Behrstock, J., Bestvina, M., and Clay, M. Growth of intersection numbers for free group automorphisms. *Journal of Topology 3*, 2 (2010), 280–310.

[4] Behrstock, J., Druţu, C., et al. Divergence, thick groups, and short conjugators. *Illinois Journal of Mathematics 58*, 4 (2014), 939–980.

[5] Behrstock, J., Druţu, C., and Mosher, L. Thick metric spaces, relative hyperbolicity, and quasi-isometric rigidity. *Mathematische Annalen 344*, 3 (2009), 543–595.

[6] Bestvina, M., Feighn, M., et al. A combination theorem for negatively curved groups. *Journal of Differential Geometry 35*, 1 (1992), 85–101.

[7] Bestvina, M., Feighn, M., and Handel, M. The Tits alternative for $Out(F_n)$ I: Dynamics of exponentially-growing automorphisms. *Annals of Math. 151* (2000), 517–623.

[8] Bestvina, M., Feighn, M., and Handel, M. The Tits alternative for $\mathrm{Out}(F_n)$. II. A Kolchin type theorem. *Ann. of Math. (2) 161*, 1 (2005), 1–59.

[9] Bestvina, M., and Handel, M. Train tracks and automorphisms of free groups. *Ann. of Math. (2) 135*, 1 (1992), 1–51.

[10] Bogopolski, O., Martino, A., Maslakova, O., and Ventura, E. The conjugacy problem is solvable in free-by-cyclic groups. *Bull. London Math. Soc. 38*, 5 (2006), 787–794.

[11] Bridson, M. R., and Groves, D. The quadratic isoperimetric inequality for mapping tori of free group automorphisms. *Mem. Amer. Math. Soc. 203*, 955 (2010), xii+152.

[12] Brinkmann, P. Hyperbolic automorphisms of free groups. *Geom. Funct. Anal. 10*, 5 (2000), 1071–1089.

[13] Button, J. Tubular free by cyclic groups and the strongest tits alternative. *arXiv preprint arXiv:1510.05842* (2015).

[14] Button, J., and Kropholler, R. Non hyperbolic free-by-cyclic and one-relator groups. *arXiv preprint arXiv:1503.01989* (2015).

[15] Cashen, C. H., and Levitt, G. Mapping tori of free group automorphisms, and the bieri–neumann–strebel invariant of graphs of groups. *Journal of Group Theory* (2014).

[16] Clay, M. l^2-torsion of free-by-cyclic groups. *arXiv preprint arXiv:1509.09258* (2015).

[17] Clay, M., and Pettet, A. Twisting out fully irreducible automorphisms. *Geometric and Functional Analysis 20*, 3 (2010), 657–689.

[18] Dison, W., and Riley, T. Hydra groups. *arXiv preprint arXiv:1002.1945* (2010).

[19] Dowdall, S., Kapovich, I., and Leininger, C. J. McMullen polynomials and lipschitz flows for free-by-cyclic groups. *To appear in J. Euro. Math. Soc.* (2013). arXiv preprint arXiv:1310.7481.

[20] Dowdall, S., Kapovich, I., and Leininger, C. J. Dynamics on free-by-cyclic groups. *Geometry & Topology 19*, 5 (2015), 2801–2899.
[21] Dowdall, S., and Taylor, S. J. The co-surface graph and the geometry of hyperbolic free group extensions. *arXiv preprint arXiv:1601.00101* (2016).
[22] Gersten, S. The automorphism group of a free group is not a CAT(0) group. *Proceedings of the American Mathematical Society 121*, 4 (1994), 999–1002.
[23] Hagen, M. F., and Przytycki, P. Cocompactly cubulated graph manifolds. *Israel Journal of Mathematics 207*, 1 (2015), 377–394.
[24] Hagen, M. F., and Wise, D. T. Cubulating hyperbolic free-by-cyclic groups: the irreducible case. *To appear in Duke Math. J.* (2013).
[25] Hagen, M. F., and Wise, D. T. Cubulating hyperbolic free-by-cyclic groups: the general case. *Geometric and Functional Analysis (GAFA) 25*, 1 (2014), 134–179.
[26] Haglund, F. Isometries of CAT(0) cube complexes are semi-simple. *arXiv:0705.3386* (2007).
[27] Haglund, F., and Paulin, F. Simplicité de groupes d'automorphismes d'espaces à courbure négative. In *The Epstein birthday schrift*. Geom. Topol., Coventry, 1998, pp. 181–248 (electronic).
[28] Hruska, G. C., and Wise, D. T. Finiteness properties of cubulated groups. *Comp. Math.*. pp. 1–58, to appear.
[29] Hsu, T., and Wise, D. T. Cubulating malnormal amalgams. *Inventiones mathematicae 199*, 2 (2015), 293–331.
[30] Kapovich, I. Algorithmic detectability of iwip automorphisms. *Bulletin of the London Mathematical Society 46*, 2 (2014), 279–290.
[31] Kapovich, I., and Lustig, M. Cannon-Thurston fibers for iwip automorphisms of F_n. *Journal of the London Mathematical Society 91*, 1 (2015), 203–224.
[32] Levitt, G. Counting growth types of automorphisms of free groups. *Geometric and Functional Analysis 19*, 4 (2009), 1119–1146.
[33] Lustig, M. *Extended Abstracts Fall 2012: Automorphisms of Free Groups*. Springer International Publishing, Cham, 2014, ch. Tree-Irreducible Automorphisms of Free Groups, pp. 67–71.
[34] Macura, N. Quadratic isoperimetric inequality for mapping tori of polynomially growing automorphisms of free groups. *Geom. Funct. Anal. (GAFA) 10*, 4 (2000), 874–901.
[35] Macura, N. Detour functions and quasi-isometries. *The Quarterly Journal of Mathematics 53*, 2 (2002), 207–239.
[36] Sageev, M. Ends of group pairs and non-positively curved cube complexes. *Proc. London Math. Soc. (3) 71*, 3 (1995), 585–617.
[37] Schleimer, S. Polynomial-time word problems. *Comment. Math. Helv. 83*, 4 (2008), 741–765.
[38] Wise, D. Cubular tubular groups. *Transactions of the American Mathematical Society 366*, 10 (2014), 5503–5521.
[39] Wise, D. T. Cubulating small cancellation groups. *GAFA, Geom. Funct. Anal. 14*, 1 (2004), 150–214.
[40] Wise, D. T. The structure of groups with a quasiconvex hierarchy. *Preprint* (2011).
[41] Woodhouse, D. J. Classifying finite dimensional cubulations of tubular groups. *arXiv preprint arXiv:1502.02619* (2015).
[42] Woodhouse, D. J. A generalized axis theorem for cube complexes. *arXiv preprint arXiv:1602.01952* (2016).

Dept. of Pure Maths and Math. Stats., University of Cambridge, Cambridge, UK
E-mail address: markfhagen@gmail.com

Dept. of Math. and Stat., McGill University, Montreal, Quebec, Canada
E-mail address: wise@math.mcgill.ca